

Improved diffusion Monte Carlo and the Brownian fan

July 13, 2012

Martin Hairer¹ and Jonathan Weare²

¹ Mathematics Department, the University of Warwick

² Mathematics Department, the University of Chicago

Email: M.Hairer@Warwick.ac.uk, Email: weare@uchicago.edu

Abstract

We propose a modification, based on the RESTART (repetitive simulation trials after reaching thresholds) and DPR (dynamics probability redistribution) rare event simulation algorithms, of the standard diffusion Monte Carlo (DMC) algorithm. The new algorithm has a lower variance per workload, regardless of the regime considered. In particular, it makes it feasible to use DMC in situations where the “naïve” generalisation of the standard algorithm would be impractical, due to an exponential explosion of its variance.

We numerically demonstrate the effectiveness of the new algorithm on a standard rare event simulation problem (probability of an unlikely transition in a Lennard-Jones cluster), as well as a high-frequency data assimilation problem. We then provide a detailed heuristic explanation of why, in the case of rare event simulation, the new algorithm is expected to converge to a limiting process as the underlying stepsize goes to 0. This is shown rigorously in the simplest possible situation of a random walk, biased by a linear potential. The resulting limiting object, which we call the “Brownian fan”, is a very natural new mathematical object of independent interest.

Keywords: Diffusion Monte Carlo, quantum Monte Carlo, rare event simulation, sequential Monte Carlo, particle filtering, Brownian fan, branching process

Contents

1	Introduction	1
2	The algorithm and a summary of our main results	4
3	Two examples	10
4	Bias and comparison of Algorithms 1 and 2	17
5	The continuous-time limit and the Brownian fan	24
6	Convergence to the Brownian fan	45

1 Introduction

Diffusion Monte Carlo (DMC) is a well established method developed within the Quantum Monte Carlo (QMC) community to compute the ground state energy (the

lowest eigenvalue) of the Hamiltonian operator

$$\mathcal{H}\psi = -\Delta\psi + V\psi$$

as well as averages with respect to the square of the corresponding eigenfunction [Kal62, GS71, And75, CA80, KL11]. It is based on the fact that the Feynman–Kac formula

$$\psi(x, t) = \mathbf{E}_x \left(f(B_t) \exp \left(- \int_0^t V(B_s) ds \right) \right), \quad (1.1)$$

connects the solution, $\psi(x, t)$, of the partial differential equation

$$\begin{aligned} \partial_t \psi &= -\mathcal{H}\psi \\ \psi(0, \cdot) &= f, \end{aligned} \quad (1.2)$$

to expectations of a standard Brownian motion B_t starting at x . For large times t , suitably normalised integrals against the solution of (1.2) approximate normalised integrals against the ground state eigenfunction of \mathcal{H} .

As we will see in the next section, DMC is an extremely flexible tool with many potential applications going well beyond quantum Monte Carlo. We will introduce a generalised version of DMC capable of solving problems arising in a wide variety of application areas. In fact, sequential importance sampling [dDG05, Liu02] is a particular application of DMC in this general framework. In that setting the goal is to compute normalised expectations of the form

$$\frac{\mathbf{E} \left(f(y_t) \exp \left(- \sum_{t_k \leq t} V(k, y_{t_k}) \right) \right)}{\mathbf{E} \exp \left(- \sum_{t_k \leq t} V(k, y_{t_k}) \right)},$$

where the $V(k, \cdot)$ are functions encoding, for example, the likelihood of sequentially arriving observations (at times $t_1 < t_2 < t_3 < \dots$) given the state of some Markov process y_t . A central component of a sequential importance sampling scheme is a so-called “resampling” step (see e.g. [dDG05, Liu02]) in which copies of a system are weighted by the ratio of two densities and then resampled to produce an unweighted set of samples. Some popular resampling algorithms are adaptations of the generalised version of DMC that we will present below in Algorithm 1 (e.g. residual resampling as described in [Liu02]). Our results suggest that sequential importance sampling schemes could be improved (in some cases dramatically) by building the resampling step from our modification of DMC in Algorithm 2.

In addition to existing algorithms such as sequential Monte Carlo that can be naturally cast as DMC, one can imagine several unexplored uses of DMC. For example, we will show that the generalisation of DMC in Algorithm 1 below could potentially be used to compute approximations to quantities of the form

$$\mathbf{E} \left(f(y_t) \exp \left(- \int_0^t V(y_t) dy_t \right) \right),$$

or

$$\mathbf{E}(f(y_t) \exp(-V(y_t))) . \quad (1.3)$$

where y_t is a diffusion. The second of these expectations is particularly interesting. We will see that, in the context of computing expectations of this form, DMC becomes a

rare event simulation technique (see e.g. [Buc04, AG10] and the references therein), i.e. a tool for efficiently generating samples of extremely low probability events and computing their probabilities. Such tools can be used to compute the probability or frequency of dramatic fluctuations in a stochastic process. Those fluctuations might, for example, characterise a reaction in a chemical system or a failure event in an electronic device (see e.g. [FS96, Buc04]).

Unfortunately, in many of the most interesting settings, the natural generalisation of DMC introduced in Algorithm 1 exhibits a dramatic instability. We show that this instability can be corrected by a slight enlargement of the state space to include a parameter dictating which part of the state space a sample is allowed to visit. The resulting method is introduced in Algorithm 2 in Section 2 below. This modification is inspired by the RESTART and DPR rare event simulation algorithms [VAVA91, HT99, DD11]. The RESTART and DPR algorithms bias an underlying Markov process by splitting the trajectory of the process into multiple trajectories and then appropriately reweighing the resulting trajectories. The splitting occurs every time the process moves from one set to another in a predefined partition of the state space. The key feature of the algorithms that we borrow is a restriction placed on the trajectories resulting from a splitting event. The new trajectories are only allowed to explore a certain region of state space and are eliminated when they exit this region. Our modification of DMC is similar, except that our branching rule is not necessarily based on a fixed partition of state space.

In Section 3, we demonstrate the stability and utility of the new method in two applications. In the first one, we use the method to compute the probability of very unlikely transitions in a small cluster of Lennard-Jones particles. In the second example we show that the method can be used to efficiently assimilate high frequency observational data from a chaotic diffusion. In addition to these computational examples, we provide an in-depth mathematical study of the new scheme (Algorithm 2). We show that regardless of parameter regime and even underlying dynamics (i.e. y_t need not be related to a diffusion) the estimator generated by the new algorithm has lower variance per workload than the DMC estimator.

In addition, by focusing on a particular asymptotic regime (small time-discretisation parameter) we are also able to rigorously establish several dramatic results concerning the stability of Algorithm 2 in settings in which the straightforward generalisation of DMC is unstable. In Section 5, we take a closer look at the situation (1.3) used for rare event simulation. We provide a heuristic explanation why one would expect our algorithm to converge to a limiting process as the discretisation parameter converges to 0 and we give a non-rigorous characterisation of the limiting object which we call the “Brownian fan.” We also provide a detailed study of the fine properties of this limiting process in the simplest case when y_t is a random walk (rescaled so that it converges to Brownian motion) and the biasing potential V is linear. The Brownian fan is interesting in its own right from a mathematical perspective and does not seem to have been studied before, though it is very closely related to the “Virgin island model” introduced in [Hut09].

The final part of this article, Section 6, is then devoted to a rigorous proof of the convergence of the output of Algorithm 2 to the Brownian fan for any sufficiently light-tailed one-step distribution of the underlying random walk. This result is formulated in Theorem 6.11. The Brownian fan has the unusual property that, while the number of particles alive at any fixed time t is finite with moments of all orders, there is a dense set of exceptional times at which it is infinite, due to the fact that offspring are created at infinite rate. As a consequence, an immediate challenge in our study of the Brownian

fan is finding a state space in which it is a well-behaved Markov process (i.e. Feller and with continuous sample paths).

1.1 Notations

For any space \mathcal{Y} , we will denote by $M_+(\mathcal{Y})$ the space of finite positive measures on \mathcal{Y} , endowed with the topology of weak convergence, together with the convergence of total mass. Given a random variable X , we denote its law by $\mathcal{D}(X)$, except in some cases where we introduce dedicated notations.

We will often use the notation $a \lesssim b$ as a shorthand for the inequality $a \leq Cb$ for some constant C . The dependence of C on other quantities will usually be clear from the context, and will be indicated when ambiguities may arise. We will also use the standard notations $a \wedge b = \min\{a, b\}$, $a \vee b = \max\{a, b\}$, and $\lfloor a \rfloor = \max\{i \in \mathbb{Z} : i \leq a\}$.

Acknowledgements

We would like to thank H. Weber for numerous discussions on this article and S. Vollmer for pointing out a number of inaccuracies found in earlier drafts. We would also like to thank J. Goodman and E. Vanden-Eijnden for helpful conversations and advice during the early stages of this work. Financial support for MH was kindly provided by EPSRC grant EP/D071593/1, by the Royal Society through a Wolfson Research Merit Award, and by the Leverhulme Trust through a Philip Leverhulme Prize. JW was supported by NSF through award DMS-1109731.

2 The algorithm and a summary of our main results

With the potential term V removed, the PDE (1.2) is simply the familiar Fokker-Planck (or forward Kolmogorov) equation and one can of course approximate (1.1) (with $V = 0$) by

$$\hat{f}_t = \frac{1}{M} \sum_{j=1}^M f(w_t^{(j)}),$$

where the $w^{(j)}$ are M independent realisations of a Brownian motion. The probabilistic interpretation of the source term $V\psi$ in the PDE is that it is responsible for the killing and creation of sample trajectories of w . In practice one cannot exactly compute the integral appearing in the exponent in (1.1). Common practice (assuming that $t_k = k\varepsilon$) is to replace the integral by the approximation

$$\int_0^t V(w_s) ds \approx \sum_{k=0}^{\lfloor t/\varepsilon \rfloor - 1} \frac{1}{2} (V(w_{t_{k+1}}) + V(w_{t_k})) \varepsilon,$$

where $\varepsilon > 0$ is a small time-discretisation parameter. DMC then approximates

$$\langle f \rangle_t = \mathbf{E} \left(f(w_t) \exp \left(- \sum_{k=0}^{\lfloor t/\varepsilon \rfloor - 1} \frac{1}{2} (V(w_{t_{k+1}}) + V(w_{t_k})) \varepsilon \right) \right). \quad (2.1)$$

The version of DMC that we state below in Algorithm 1 is a slight generalisation of the standard algorithm and approximates

$$\langle f \rangle_t = \mathbf{E} \left(f(y_t) \exp \left(- \sum_{k=0}^{\lfloor t/\varepsilon \rfloor - 1} \chi(y_{k\varepsilon}, y_{(k+1)\varepsilon}) \right) \right), \quad (2.2)$$

where y_t is any Markov process and χ is any function of two variables. Because it will be the case in our examples as well as in our in-depth case study, we have assumed that the Markov process is sampled at equal time intervals of length ε . For simplicity we will henceforth also assume that t is a multiple of ε . These assumptions are by no means necessary. The DMC algorithm proceeds as follows:

Algorithm 1 *Slightly generalised DMC*

1. Begin with M copies $x_0^{(j)} = x_0$ and $k = 0$.
2. At step k there are $N_{k\varepsilon}$ samples $x_{k\varepsilon}^{(j)}$. Evolve each of these ε units of time under the underlying dynamics to generate $N_{k\varepsilon}$ values

$$\tilde{x}_{(k+1)\varepsilon}^{(j)} \sim \mathbf{P}(y_{(k+1)\varepsilon} \in dx \mid y_{k\varepsilon} = x_{k\varepsilon}^{(j)}).$$

3. For each $j = 1, \dots, N_{k\varepsilon}$, let

$$P^{(j)} = e^{-\chi(x_{k\varepsilon}^{(j)}, \tilde{x}_{(k+1)\varepsilon}^{(j)})}$$

and set

$$N^{(j)} = \lfloor P^{(j)} + u^{(j)} \rfloor,$$

where the $u^{(j)}$ are independent $\mathcal{U}(0,1)$ random variables.

4. For $j = 1, \dots, N_s$, and for $i = 1, \dots, N^{(j)}$ set

$$x_{(k+1)\varepsilon}^{(j,i)} = \tilde{x}_{(k+1)\varepsilon}^{(j)}.$$

5. Finally, set $N_{(k+1)\varepsilon} = \sum_{j=1}^{N_{k\varepsilon}} N^{(j)}$ and list the $N_{(k+1)\varepsilon}$ vectors $\{x_{(k+1)\varepsilon}^{(j,i)}\}$ as $\{x_{(k+1)\varepsilon}^{(j)}\}_{j=1}^{N_{(k+1)\varepsilon}}$.

6. At time t , produce the estimate

$$\hat{f}_t = \frac{1}{M} \sum_{j=1}^{N_t} f(x_t^{(j)}).$$

Here, the notation $\mathcal{U}(a, b)$ was used to refer to the law of a random variable that is uniformly distributed in the interval (a, b) . Below we will refer to the $x_t^{(j)}$ as *particles* and to a collection of particles as an *ensemble*. Algorithm 1 results in an unbiased estimate of (2.2), i.e.

$$\mathbf{E}^1 \hat{f}_t = \langle f \rangle_t.$$

For reasonable choices of χ and f (e.g. $\sup_{k \leq t/\varepsilon} \langle 1 \rangle_{k\varepsilon} < \infty$ and $\langle |f| \rangle_t < \infty$) the law of large numbers then implies that

$$\lim_{M \rightarrow \infty} \hat{f}_t = \langle f \rangle_t.$$

We use the symbol \mathbf{E}^1 for expectations under the rules of Algorithm 1 to distinguish them for expectations under the rules of our modified algorithm (Algorithm 2 below) which will simply be denoted by \mathbf{E} .

Of course if we set

$$\chi(x, y) = \frac{1}{2}(V(x) + V(y))\varepsilon, \quad (2.3)$$

then we are back to the quantum Monte Carlo setting. We might, however, choose

$$\chi(x, y) = V(x)(y - x), \quad (2.4)$$

which, if y_t approximates a diffusion (which we also denote by y_t), formally corresponds to an approximation of

$$\langle f \rangle_t = \mathbf{E} \left(f(y_t) \exp \left(- \int_0^t V(y_s) dy_s \right) \right),$$

or

$$\chi(x, y) = V(y) - V(x), \quad (2.5)$$

corresponding to an approximation of

$$\langle f \rangle_t = e^{V(x_0)} \mathbf{E} \left(f(y_t) e^{-V(y_t)} \right).$$

The generalisation of DMC resulting from various choices of χ is not frivolous. In Section 3 we will see that, with χ of form (2.5), one can design potentially very useful rare event and particle filtering schemes. Unfortunately, when y_t is a diffusion or a discretisation of a diffusion, the resulting Algorithms behave extremely poorly in the small discretisation size limit ($\varepsilon \rightarrow 0$). The reason is simple: on a time interval ε , a typical step of Brownian motion moves by $\mathcal{O}(\sqrt{\varepsilon})$. Therefore, unlike for (2.3), for both (2.4) and (2.5), the typical size of $\chi(x, y)$ is of order $\sqrt{\varepsilon}$. As a consequence, at each step, the probability that a particle either dies or spawns a child is itself of order $\sqrt{\varepsilon}$. Without modification, as $\varepsilon \rightarrow 0$, the fluctuations in the number of particles, $N_{[t/\varepsilon]}$, will grow wildly and the process will die out before a time of order 1 with higher and higher probability, thus causing the variance of our estimator to explode. This phenomenon is already evident in the simple case of a Brownian motion,


$$y_{(k+1)\varepsilon} = y_{k\varepsilon} + \sqrt{\varepsilon} \xi_{k+1}, \quad X(0) = 0, \quad (2.6)$$

with

$$\chi(x, y) = y - x \quad (2.7)$$

(a choice consistent with either (2.4) or (2.5)). Here the ξ_k are independent mean 0 and variance 1 Gaussian random variables. We will analyse this example (with more general ξ_k) in great depth in Sections 5 and 6. For the moment, Figure 2 demonstrates the dramatic failure of the straightforward generalisation of DMC in Algorithm 1 on this simple example. There we plot the logarithm of the second moment of the number of particles in the ensemble at time 1 (i.e. after $1/\varepsilon$ steps of the algorithm) versus $-\log \varepsilon$ for several values of ε . Clearly, $\mathbf{E}(N_1^2)$ is growing exponentially in $1/\varepsilon$. This instability can be removed by carrying out the branching steps (Steps 3-5) only every $\mathcal{O}(1)$ units of time instead of every ε units of time and accumulating weights for the particles in the intervening time intervals. However, this small ε regime cleanly highlights a serious (and unnecessary) deficiency in DMC. In contrast to Algorithm 1, the method that we will describe in Algorithm 2 below, appears to be stable as ε vanishes (a property that will be confirmed later in Proposition 6.2).

Our modification of DMC described in Algorithm 2 below not only suppresses the small ε instability just described, but results in a more robust and efficient algorithm in any context. Inspired by the RESTART and DPR algorithms studied in [VAVA91, HT99, DD11] we append to each particle a “ticket” θ limiting the region that the particle is

Figure 1: Second moment of N versus stepsize.

allowed to visit. Samples are eliminated when and only when they violate their ticket values. As a consequence, particles typically survive for a much longer time, and in many situations there is a well-defined limiting algorithm when $\varepsilon \rightarrow 0$. More precisely, the modified algorithm proceeds as follows:

Algorithm 2 *Ticketed DMC*

1. Begin with M copies $x_0^{(j)} = x_0$. For each $j = 1, \dots, M$ choose an independent random variable $\theta_0^{(j)} \sim \mathcal{U}(0, 1)$.
2. At step k there are $N_{k\varepsilon}$ samples $(x_{k\varepsilon}^{(j)}, \theta_{k\varepsilon}^{(j)})$. Evolve each of the $x_{k\varepsilon}^{(j)}$ one step to generate $N_{k\varepsilon}$ values

$$\tilde{x}_{(k+1)\varepsilon}^{(j)} \sim \mathbf{P}(y_{(k+1)\varepsilon} \in dx \mid y_{k\varepsilon} = x_{k\varepsilon}^{(j)}).$$

3. For each $j = 1, \dots, N_{k\varepsilon}$, let

$$P^{(j)} = e^{-\chi(x_{k\varepsilon}^{(j)}, \tilde{x}_{(k+1)\varepsilon}^{(j)})}. \quad (2.8)$$

If $P^{(j)} < \theta_{k\varepsilon}^{(j)}$ then set

$$N^{(j)} = 0.$$

If $P^{(j)} \geq \theta_{k\varepsilon}^{(j)}$ then let

and set

$$N^{(j)} = \lfloor P^{(j)} + u^{(j)} \rfloor \vee 1, \quad (2.9)$$

where $u^{(j)}$ are independent $\mathcal{U}(0, 1)$ random variables.

4. For $j = 1, \dots, N_{k\varepsilon}$, if $N^{(j)} > 0$ set

$$x_{(k+1)\varepsilon}^{(j,1)} = \tilde{x}_{(k+1)\varepsilon}^{(j)} \quad \text{and} \quad \theta_{(k+1)\varepsilon}^{(j,1)} = \frac{\theta_{k\varepsilon}^{(j)}}{P^{(j)}}$$

and for $i = 2, \dots, N^{(j)}$

$$x_{(k+1)\varepsilon}^{(j,i)} = \tilde{x}_{(k+1)\varepsilon}^{(j)} \quad \text{and} \quad \theta_{(k+1)\varepsilon}^{(j,i)} \sim \mathcal{U}((P^{(j)})^{-1}, 1).$$

5. Finally set $N_{(k+1)\varepsilon} = \sum_{j=1}^{N_{k\varepsilon}} N^{(j)}$ and list the $N_{(k+1)\varepsilon}$ vectors $\{x_{(k+1)\varepsilon}^{(j,i)}\}$ as $\{x_{(k+1)\varepsilon}^{(j)}\}_{j=1}^{N_{(k+1)\varepsilon}}$.

6. At time t produce the estimate

$$\hat{f}_t = \frac{1}{M} \sum_{j=1}^{N_t} f(x_t^{(j)}) .$$

Remark 2.1 Both here and in Algorithm 1, we could have replaced the random variable $[P^{(j)} + u^{(j)}]$ by any integer-valued random variable with mean $P^{(j)}$. The choice given here is natural, since it is the one that minimises the variance. From a mathematical perspective, the analysis would have been slightly easier if we chose instead to use Poisson distributed random variables.

We should first establish that this new algorithm “does the job” in the sense that if the steps in Algorithm 1 are replaced by those in Algorithm 2, the resulting ensemble of particles still produces an unbiased estimate of the quantity $\langle f \rangle_t$ defined in (2.2). This is the subject of Theorem 4.1 below which establishes that, indeed, for any bounded test function f and any choice of χ ,

$$\mathbf{E} \hat{f}_t = \langle f \rangle_t . \quad (2.10)$$

We have already mentioned that a considerable portion of this paper is devoted to showing that our modification of DMC, Algorithm 2, does not suffer the small ε instability observed in Algorithm 1. In that asymptotic regime Algorithm 2 dramatically outperforms the straightforward generalisation of DMC in Algorithm 1. However, a surprising side-effect of our modification is that Algorithm 2 is superior to Algorithm 1 in all contexts. Before describing our theoretical results in the small ε limit, we will therefore compare Algorithms 1 and 2 independently of a specific regime (discretisation stepsize, choice of χ , etc).

The key tool in carrying out this comparison is Lemma 4.4 in Section 4. That lemma establishes that by simply randomising all tickets uniformly between 0 and 1 at each step of Algorithm 2, one obtains a process that, in law, is identical to the one generated by Algorithm 1. With this observation in hand we are able to establish in Theorem 4.3 that the estimator produced by Algorithm 2 always has lower variance than Algorithm 1, i.e. we prove that

$$\mathbf{var} \hat{f}_t \leq \mathbf{var}^1 \hat{f}_t , \quad (2.11)$$

holds for any bounded test function f , any underlying Markov chain Y , and any choice of χ . In expression (2.11) we have again distinguished expectations under the rules of Algorithm 1 from all other expectations by a superscript 1.

In comparing Algorithms 1 and 2, it is not enough to simply compare the variances of the corresponding estimators: one should also compare their respective costs. Since the dominant contribution to any cost difference in the two algorithms will come from differences in the number of times one must evolve a particle from one time step to the next we compare the expectations of *workload*

$$\mathcal{W}_t = \sum_{k=0}^{t/\varepsilon-1} N_{k\varepsilon}$$

under the two rules. However, by (2.10) with $f \equiv 1$ we see that

$$\mathbf{E}^1 \mathcal{W}_t = \mathbf{E} \mathcal{W}_t = M \sum_{k=0}^{t/\varepsilon-1} \mathbf{E} \exp\left(-\sum_{j=0}^{k-1} \chi(y_{j\varepsilon}, y_{(j+1)\varepsilon})\right),$$

so that the expected cost of both algorithms is the same. In fact, the proof of Theorem 4.3 can be modified slightly to show that the variance of the workload of Algorithm 2 is always lower than the variance of workload of Algorithm 1. These remarks leave little room for ambiguity. Algorithm 2 is more efficient than Algorithm 1 in any setting.

Let us return now to the behaviour of the process generated by Algorithm 2 in the $\varepsilon \rightarrow 0$ limit in the special case delineated by equations (2.6) and (2.7) (but allowing more generality in the ξ_k variables). This topic is the subject of Sections 5 and 6. We begin in Section 5 by motivating the construction of the continuous-time limit of the process generated by Algorithm 2 in the special case in which Y is an appropriately scaled random walk and $\chi(x, y) = y - x$. The process that we construct there in Definition 5.8 is built recursively by successive realisations of a Poisson point process in a space of excursions of y_t . We call this process the *Brownian fan*. Loosely speaking, the Brownian fan is a branching process obtained in the following way. Start with a (usually finite) number of initial particles on \mathbf{R} that furthermore come each with a tag $v \in \mathbf{R}$, such that the position x satisfies $x > v$. Each of these particles, which we call the ancestor particles, perform independent Brownian motions until they hit the barrier $x = v$, where they get killed. Each of these particles independently produces offspring according to the following mechanism. Denoting by \mathcal{E} the space of excursions on \mathbf{R} , let M be the Poisson point process on $\mathbf{R} \times \mathcal{E}$ with intensity measure $a dt \otimes Q$ for some $a > 0$, where Q is Itô's Brownian excursion measure [Itô72]. If (t, w) is one of the points of M and the corresponding ancestor particle is alive at time t and located at x_t , then it gives raise to an offspring which performs the motion described by w , translated so that the origin of the excursion lies at (t, x_t) . This mechanism is then repeated recursively for each new particle created in this way.

The Brownian fan has a number of very nice properties. For example, as established by Theorem 5.13, the continuous-time analogue of the number of particles at time t , N_t , corresponding to the Brownian fan satisfies

$$\sup_{t>0} \mathbf{E} \exp(\lambda_t N_t) < \infty \quad (2.12)$$

for some continuous decreasing function $\lambda_t > 0$. As established in Proposition 5.15, the bound in (2.12) implies that the continuous-time analogue of the workload,

$$\mathcal{W}_t = \int_0^t N_s ds$$

is very nearly a differentiable function of time,

$$\sup_{t \leq T} \lim_{h \rightarrow 0} \frac{|\mathcal{W}_{t+h} - \mathcal{W}_t|}{h |\log h|} < \infty.$$

This is close to optimal, since there turns out to be a dense set of exceptional times at which there are infinitely many particles alive (see the argument given at the end of Section 5.2.3 below).

Finally, in Proposition 5.18 we establish that the Brownian fan is a Feller process in a suitable state space which, when combined with the continuity property established in Proposition 6.8 implies that the Brownian fan is a strong Markov process.

The usual method by which one attempts to characterise the continuous-time limit of a sequence of discrete-time processes involves studying the limit of the generators of the discrete-time processes. In our case, as described in Section 5.3.4 the discrete-time generators do *not* converge to the correct limit. Surprisingly (and rather confusingly), they do actually converge to the generator of a Brownian fan, but unfortunately with the wrong parameters! Nevertheless, without resort to generators, in Theorem 6.11 we are able to establish the convergence of the process generated by Algorithm 2 to the Brownian fan. En route to this result we establish a number of important and extremely encouraging results about the behaviour of the process generated by Algorithm 2 for finite ε . For example, in Proposition 6.2 we obtain that

$$\sup_{\varepsilon < \varepsilon_0} \mathbf{E}|N_t|^p < \infty ,$$

for any $p \geq 0$ and $\varepsilon_0 > 0$ sufficiently small where here N_t refers to the number of particles generated in Algorithm 2 and not to the Brownian fan (for which we have the even stronger result in (2.12)). In Corollary 6.10 we prove a uniform (in ε) form of continuity of the processes.

The existence of a continuous-time limit along with the many other robust features of Algorithm 2 established in the results just mentioned strongly suggest that Algorithm 2 has significant potential as a flexible and efficient computational tool. In the next section we provide numerical evidence to that effect.

3 Two examples

In the following subsections we apply our modified DMC algorithm to two simple problems. In Section 3.1 we consider the approximation of a small probability related to the escape of a diffusion from a deep potential well. We will use Algorithm 2 to bias simulation of the diffusion so that an otherwise rare event (escape from the potential well) becomes common (and then reweight appropriately). We will see that extremely low probability events can be computed with reasonably low relative error. In Section 3.2 we demonstrate the performance of Algorithm 2 in one of DMCs most common application contexts, particle filtering. We will reconstruct a hidden trajectory of a diffusion, given noisy but very frequent observations. Our comparison shows that the reconstruction produced by the modified method is dramatically more accurate than the reconstruction produced by straightforward generalisation of DMC.

3.1 Rare event simulation

In Quantum Monte Carlo applications, the function χ is (for the most part) specified by the potential V as in (2.3). In other applications however, there may be more freedom in the choice of χ to achieve some specific goal. As already mentioned earlier, the choice $\chi(x, y) = V(y) - V(x)$ turns Algorithm 2 into an estimator of

$$\langle f \rangle_t = e^{V(x_0)} \mathbf{E}(f(y_t) e^{-V(y_t)}) .$$

The design of a “good” choice of V will be discussed below. By redefining f we find that $e^{-V(x_0)} \widehat{e^V f}_t$ is an unbiased estimate of $\mathbf{E}(f(y_t))$, i.e.

$$e^{-V(x_0)} \mathbf{E} \sum_{j=1}^{N_t} e^{V(x_t^{(j)})} f(x_t^{(j)}) = \mathbf{E} f(y_t) ,$$




Figure 2: **(Left)** The initial condition x^A to generate the results in Table 1. **(Right)** A typical sample of the event $y_2 \in B$.

where the samples $x_t^{(j)}$ are generated as in Algorithm 2. This suggests choosing V to minimise the variance of $\widehat{e^V f}_t$. Intuitively this involves choosing V to be smaller in regions where f is significant. The branching step in Algorithm 2 will create more copies when V decreases, focusing computational resources in regions where V is small.

As an example, consider the overdamped Langevin dynamics

$$dy_t = -\nabla U(y_t)dt + \sqrt{2\gamma}dB_t$$

where y represents the positions of seven two-dimensional particles (i.e. $y \in \mathbb{R}^{14}$) and

$$U(x) = \sum_{i < j} 4(\|x_i - x_j\|^{-12} - \|x_i - x_j\|^{-6}).$$

In this formula x_i represents the position of the i th particle (i.e. $x_i \in \mathbb{R}^2$). In various forms this Lennard-Jones cluster is a standard non-trivial test problem in rare event simulation (see [DBC98]).

After discretising this process by the Euler scheme

$$y_{(k+1)\varepsilon} = y_{k\varepsilon} - \nabla U(y_{k\varepsilon})\varepsilon + \sqrt{2\gamma\varepsilon}\xi_{k+1}$$

with a stepsize of $\varepsilon = 10^{-3}$ we apply Algorithm 2 with

$$\chi(x, y) = V(y) - V(x), \quad V(x) = \frac{\lambda}{kT} \min_{i \geq 2} \left\{ \left\| x_i - \frac{1}{7} \sum_j x_j \right\| \right\}.$$

The properties of the Euler discretisation in the context of rare event simulation are discussed in [VEWb].

Our goal is to compute

$$\mathbf{P}_{x^A}(y_2 \in B), \quad B = \{x : V(x) < 0.1\},$$

where our initial configuration is given by $x_1^A = (0, 0)$ and for $j = 2, 3, \dots, 7$,

$$x_j^A = \left(\cos \frac{j\pi}{3}, \sin \frac{j\pi}{3} \right)$$

(see Figure 2). Starting from initial configuration $x_0 = x^A$, the particle initially at the centre of the cluster (x_1^A) will typically remain there for a length of time that increases exponentially in γ^{-1} . We are roughly computing the probability that, in 2 units of time,

γ	λ	estimate	workload	$\frac{1}{2} \left(\frac{\text{variance}}{\times \text{workload}} \right)$	brute force variance
0.4	1.9	1.125×10^{-2}	12.90	4.732×10^{-3}	1.112×10^{-2}
0.2	1.3	2.340×10^{-3}	11.64	2.344×10^{-4}	2.335×10^{-3}
0.1	1	7.723×10^{-5}	13.87	7.473×10^{-7}	7.722×10^{-5}
0.05	0.85	9.290×10^{-8}	30.84	1.002×10^{-11}	9.290×10^{-8}
0.025	0.8	1.129×10^{-13}	204.8	1.311×10^{-21}	1.129×10^{-13}

Table 1: Performance of Algorithm 2 on the Lennard-Jones cluster problem described in Section 3.1 at different temperatures (γ). A measure of the efficiency improvement of Algorithm 2 over straightforward simulation is obtained by taking the ratio of the last two columns

the particle at the centre of the cluster exchanges positions with one of the particles in the outer shell. This probability decreases exponentially in γ^{-1} as $\gamma \rightarrow 0$. In this case

$$f(x) = \mathbf{1}_B(x),$$

where $\mathbf{1}_B$ is the indicator function of B . In our simulations, λ is chosen so that the expected number of particles ending in B , $\mathbf{E}\hat{f}_2$, is (very) roughly 1. The results for several values of the temperature γ are displayed in Table 1.

The workload referenced in Table 1 is the (scaled) expected total number of ∇U evaluations per sample, i.e. the expectation of

$$\varepsilon \mathcal{W}_2 = \varepsilon \sum_{k=1}^{2/\varepsilon} N_{k\varepsilon}.$$

Note that when $V \equiv 0$, Algorithm 2 reduces to straightforward (brute force) simulation and the workload is exactly 2. Increasing λ results in the creation of more particles. This has the competing effects of increasing the workload on the one hand but reducing the variance of the resulting estimator on the other. If we let σ^2 denote the variance of one sample run of Algorithm 2 (i.e. $\mathbf{var} e^{-V(x^A)} \widehat{e^V f_2}$ with $M = 1$) and σ_{bf}^2 the variance of the random variable $f(y_2)$, then the number of samples required to compute an estimate with a statistical error of err is $M = \sigma^2/err$ and $M_{bf} = \sigma_{bf}^2/err$ for Algorithm 2 and brute force respectively. Taking account of the expected cost (scaled by ε) per sample of Algorithm 2 (reported in the “workload” column) and of brute force simulation (identically 2) we can obtain a comparison of the cost of Algorithm 2 and brute force by comparing the brute force variance to the product of the variance of the estimate produced by Algorithm 2 and one-half the workload reported in the fourth column of the table. These quantities are reported in the last two columns of the table. By taking the ratio of the two columns we obtain a measure of the relative cost per degree of accuracy of the two methods. One can see that the speedup provided by Algorithm 2 becomes dramatic for smaller values of γ . The variance of the brute force estimator is computed from the values in the third column by

$$\text{brute force variance} = \mathbf{P}_{x^A}(y_2 \in B)(1 - \mathbf{P}_{x^A}(y_2 \in B)).$$

Only the estimates in the first two rows of Table 1 were compared against straightforward simulation. Those estimates agreed to the appropriate precision.

The choice of V used in the above simulations is certainly far from optimal. Given the similarities in this rare event context between Algorithm 2 and the DPR algorithm, it

should be straightforward to adapt the results of Dean and Dupuis in [DD11] identifying optimal choices of V in the small γ limit. In fact, we suspect (but do not prove) that the optimal choice of V at any temperature is given by the time-dependent function

$$V(k, x) = -\log \mathbf{P}(y_2 \in B \mid y_{k\varepsilon} = x).$$

This choice is clearly impractical as it requires knowing the probability that we are trying to compute. Even asymptotic estimates based on taking an appropriate low γ limit of this choice are often not practical so it is worth noting that, even with a rather clumsy choice of V , our scheme is still able to produce impressive results to fairly low temperatures. We fully expect however that without a very carefully chosen V , the cost of achieving a fixed relative accuracy for our scheme will grow exponentially with γ^{-1} as $\gamma \rightarrow 0$. This characteristic is common in rare event simulation methods.

3.2 Non-linear filtering

In this section our goal is to reconstruct a sample path of the solution to the stochastic differential equation

$$dy_t^{(1)} = 10(y_t^{(1)} - y_t^{(2)}) dt + \sqrt{2} dB_t^{(1)}, \quad (3.1)$$

$$dy_t^{(2)} = (y_t^{(1)}(28 - y_t^{(3)}) - y_t^{(2)}) dt + \sqrt{2} dB_t^{(2)}, \quad (3.2)$$

$$dy_t^{(3)} = \left(y_t^{(1)} y_t^{(2)} - \frac{8}{3} y_t^{(3)} \right) dt + \sqrt{2} dB_t^{(3)}, \quad (3.3)$$

with

$$y_0 = \begin{pmatrix} -5.91652 \\ -5.52332 \\ 24.57231 \end{pmatrix}.$$

The deterministic version of this system is the famous Lorenz 63 chaotic ordinary differential equation [Lor63]. The stochastic version above is commonly used in tests of on-line filtering strategies (see e.g. [MGG94, MTAC12]).

The path of y_t is revealed only through the 3-dimensional noisy observation process

$$dh_t = y_t dt + 0.1 d\tilde{B}_t \quad (3.4)$$

where the 3-dimensional Brownian motion \tilde{B}_t is independent of B_t . Our task is to reconstruct the path of y_t given a realisation of h_t . More precisely, our goal is to sample from (and compute moments of) the conditional distribution of y_t given \mathcal{F}_t^h , where \mathcal{F}^h is the filtration generated by h . One can verify that expectations with respect to this conditional distribution can be written as

$$\frac{\mathbf{E}^B \left(f(y_t) \exp \left(-10 \int_0^t \langle y_s, dh(s) \rangle - 50 \int_0^t \|y_s\|^2 ds \right) \right)}{\mathbf{E}^B \exp \left(-10 \int_0^t \langle y_s, dh_s \rangle - 50 \int_0^t \|y_s\|^2 ds \right)}, \quad (3.5)$$

where the superscript B on the expectations indicates that they are expectations over B_t only and not over \tilde{B}_t , i.e. the trajectory of h_t is fixed. We will focus on estimating the mean of this distribution ($f(x) = x$) which we will refer to as the reconstruction of the hidden signal.

As always, we should first discretise the problem. The simplest choice is to again replace (3.1) by its Euler discretisation with parameter ε . The observation process (3.4) can be replaced by

$$h_{(k+1)\varepsilon} - h_{k\varepsilon} \approx y_{k\varepsilon} \varepsilon + 0.1 \sqrt{\varepsilon} \eta_{k+1} \quad (3.6)$$




Figure 3: Hidden signal (black), together with the observed signal in gray.

where the η_k are independent Gaussian random variables with mean 0 and identity covariance. We again set $\varepsilon = 10^{-4}$. With this choice for h_t , it is equivalent to assume that the observation process is the sequence of increments $h_{(k+1)\varepsilon} - h_{k\varepsilon}$ instead of H itself. To emphasise that we will be conditioning on the values of $h_{(k+1)\varepsilon} - h_{k\varepsilon}$ and not computing expectations with respect to these variables we will use the notation Δ_k to denote a specific realisation of $h_{(k+1)\varepsilon} - h_{k\varepsilon}$. In Figure 3.2 we plot a trajectory of the resulting discrete process y and observations $h_{(k+1)\varepsilon} - h_{k\varepsilon}$.

Note that given a particular state $y_{k\varepsilon} = x$, the probability of observing $h_{(k+1)\varepsilon} - h_{k\varepsilon} = \Delta$ is

$$\exp\left(-\frac{\|x\varepsilon - \Delta\|^2}{0.02\varepsilon}\right)$$

and expectations of the discretised process y given the $h_{(k+1)\varepsilon} - h_{k\varepsilon} = \Delta_k$ observations can be computed by the formula

$$Z_k^{-1} \mathbf{E}\left(y_{k\varepsilon} \exp\left(-\sum_{j=1}^k \frac{\|y_{j\varepsilon} \varepsilon - \Delta_j\|^2}{0.02\varepsilon}\right)\right), \quad (3.7)$$

where the normalisation constant Z_k is given by

$$Z_k = \mathbf{E} \exp\left(-\sum_{j=1}^k \frac{\|y_{j\varepsilon} \varepsilon - \Delta_j\|^2}{0.02\varepsilon}\right).$$

In the small ε limit, formula (3.7) (with $k = t/\varepsilon$) indeed converges to (3.5) (see [Cri11]).

Expression (3.7) suggests applying Algorithms 1 or 2 with

$$\chi(x, y) = \frac{\|y\varepsilon - \Delta_{k+1}\|^2}{0.02\varepsilon}$$

at each time step. Below it will be convenient to consider the resulting choice of $P^{(i)}$,

$$P^{(i)} = \exp\left(-\frac{\|\tilde{x}_{(k+1)\varepsilon}^{(i)} - \Delta_{k+1}\|^2}{0.02\varepsilon}\right)$$

in Algorithm 2. With this choice, the expected number of particles at the k th step would be

$$M \mathbf{E} \exp \left(- \sum_{j=1}^k \frac{\|y_{j\varepsilon} \varepsilon - \Delta_j\|^2}{0.02\varepsilon} \right),$$

a quantity that may decay very rapidly with k . Since our goal is to compute conditional expectations we are free to normalise $P^{(i)}$ by any constant. A tempting choice is

$$P^{(i)} = \frac{Z_k}{Z_{k+1}} \exp \left(- \frac{\|\tilde{x}_{(k+1)\varepsilon}^{(i)} \varepsilon - \Delta_{k+1}\|^2}{0.02\varepsilon} \right)$$

which would result in $\mathbf{E}N_{k\varepsilon} = M$ for all k . Unfortunately we do not know the conditional expectation in the denominator of this expression. However, for a large number of particles $N_{k\varepsilon}$ we have the approximation

$$\begin{aligned} \frac{1}{N_{k\varepsilon}} \sum_{i=1}^{N_{k\varepsilon}} \exp \left(- \frac{\|\tilde{x}_{(k+1)\varepsilon}^{(i)} \varepsilon - \Delta_{k+1}\|^2}{0.02\varepsilon} \right) \\ \approx \mathbf{E} \left(\exp \left(- \frac{\|y_{(k+1)\varepsilon} \varepsilon - \Delta_{k+1}\|^2}{0.02\varepsilon} \right) \mid \Delta_1, \dots, \Delta_k \right) = \frac{Z_{k+1}}{Z_k}. \end{aligned}$$

This suggests using

$$P^{(i)} = \frac{\exp \left(- \frac{\|\tilde{x}_{(k+1)\varepsilon}^{(i)} \varepsilon - \Delta_{k+1}\|^2}{0.02\varepsilon} \right)}{\frac{1}{N_{k\varepsilon}} \sum_{l=1}^{N_{k\varepsilon}} \exp \left(- \frac{\|\tilde{x}_{(k+1)\varepsilon}^{(l)} \varepsilon - \Delta_{k+1}\|^2}{0.02\varepsilon} \right)}, \quad (3.8)$$

which will guaranty that $\mathbf{E}N_{k\varepsilon} = M$ for all k .

In fact, in practical applications it is important to have even more control over the population size. If we use $P^{(i)}$ in formula (3.8) in our DMC algorithms, the resulting N -process will be a Martingale bounded below by the absorbing state $N = 0$. As such it will eventually reach that absorbing state, an inconvenient predicament in practical implementations. Stronger control on the number of particles can be achieved in many ways (e.g. by resampling strategies [dDG05, Liu02]). We choose to make the simple modification

$$P^{(i)} = M \frac{\exp \left(- \frac{\|\tilde{x}_{(k+1)\varepsilon}^{(i)} \varepsilon - \Delta_{k+1}\|^2}{0.02\varepsilon} \right)}{\sum_{l=1}^{N_{k\varepsilon}} \exp \left(- \frac{\|\tilde{x}_{(k+1)\varepsilon}^{(l)} \varepsilon - \Delta_{k+1}\|^2}{0.02\varepsilon} \right)}, \quad (3.9)$$

which results in an expected number of particles at step $k + 1$ of M independently of the details of the step k ensemble. Formula (3.9) depends on the entire ensemble of walkers and does not fit strictly within the confines of Algorithms 1 and 2. This choice will lead to estimators with a small, M -dependent, bias. All sequential Monte Carlo strategies with ensemble population control suffer a similar bias.

The results of a test of Algorithm 2 with this choice of $P^{(i)}$ are presented in Figure 5 for $M = 10$. The true trajectory of y is drawn as a dotted line, while our reconstruction is drawn as a solid line. Note that the dotted line is nearly completely hidden behind the solid line, indicating an accurate reconstruction. In Figure 4 we show the results of the same test with Algorithm 2 replaced by Algorithm 1. The resulting method is




Figure 4: Reconstruction of the three components of the signal using Algorithm 1. The hidden signal is shown as a dotted line.




Figure 5: Reconstruction of the three components of the signal using Algorithm 2. The hidden signal is shown as a dotted line.

very similar to a standard particle filter with the common residual resampling step (see [Liu02]).

With our small choice of ε , one might expect that the number of particles generated by Algorithm 1 would explode or vanish almost instantly. Our choice of $P^{(i)}$ prevents large fluctuations in N . We expect instead that the number of truly distinguishable particles in the ensemble (in the sense that they are not very nearly in the same positions) will drop to 1 instantly, leading to a very low resolution scheme and a poor reconstruction. This is supported by the results shown in Figure 4.

The fact that we obtain such a good reconstruction in Figure 5 with only 10 particles indicates that this is not a particularly challenging filtering problem. Challenging filtering problems and more involved methods for dealing with them are discussed in [VEWa]. The purpose of this example is to demonstrate that by simply removing unnecessary variance from a particle filter (via Algorithm 2) one can improve performance. This improvement is illustrated dramatically in our small ε setting. In practice it is advisable to carry out the copying and killing steps in Algorithm 2 less frequently, accumulating particle weights in the intervening time intervals. Nonetheless, any variance in the resulting estimate will be unnecessary amplified if one uses a method based on Algorithm 1 instead of Algorithm 2.

4 Bias and comparison of Algorithms 1 and 2

In this section we establish several general properties of Algorithms 1 and 2. In contrast to later sections we do not focus here on any specific choice of the Markov process y or the function χ . For simplicity in our displays we will assume in this section that y is time-homogenous, i.e. that its transition probability distribution is independent of time. As above we use \mathbf{E}^1 and \mathbf{E} to denote expectations under the rules of Algorithms 1 and 2 respectively. In the proof of each result we will assume that $M = 1$ in Algorithms 1 and 2. The results follow for $M > 1$ by the independence of the processes generated from each copy of the initial configuration.

Our first main result in this section establishes that Algorithms 1 and 2 are unbiased in the following sense.

Theorem 4.1 *Algorithm 2 produces an unbiased estimator, namely*

$$\mathbf{E} \hat{f}_t = \mathbf{E} \left(f(y_t) \exp \left(- \sum_{k=0}^{t/\varepsilon-1} \chi(y_{k\varepsilon}, y_{(k+1)\varepsilon}) \right) \right). \quad (4.1)$$

The expression also holds with \mathbf{E} replaced by \mathbf{E}^1 .

The proof of Theorem 4.1 is similar to the proof of Theorem 3.3 in [DD11] and requires the following lemma.

Lemma 4.2 *For any bounded, non-negative function F on $\mathbb{R}^d \times \mathbb{R}$ we have*

$$\mathbf{E} \left(\sum_{j=1}^{N_\varepsilon} F(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) \mid x_0, \tilde{x}_\varepsilon^{(1)} \right) = e^{-\chi(x_0, \tilde{x}_\varepsilon^{(1)})} \int_0^1 F(\tilde{x}_\varepsilon^{(1)}, u) du. \quad (4.2)$$

The expression also holds with \mathbf{E} replaced by \mathbf{E}^1 .

Proof. From the description of the algorithm and, in particular the fact that given x_0 , and $\tilde{x}_\varepsilon^{(1)}$, the tickets are all independent of each other, we obtain

$$\begin{aligned} \mathbf{E} \left(\sum_{j=1}^{N_\varepsilon} F(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) \middle| x_0, \tilde{x}_\varepsilon^{(1)} \right) &= \int_0^1 F(\tilde{x}_\varepsilon^{(1)}, u e^{\chi(x_0, \tilde{x}_\varepsilon^{(1)})}) \mathbf{1}_{(e^{-\chi(x_0, \tilde{x}_\varepsilon^{(1)})} \geq u)} du \\ &\quad + \frac{e^{-\chi(x_0, \tilde{x}_\varepsilon^{(1)})} - 1}{1 - e^{\chi(x_0, \tilde{x}_\varepsilon^{(1)})}} \int_{e^{\chi(x_0, \tilde{x}_\varepsilon^{(1)})}}^1 F(\tilde{x}_\varepsilon^{(1)}, u) du \mathbf{1}_{(\chi(x_0, \tilde{x}_\varepsilon^{(1)}) \leq 0)}. \end{aligned}$$

The first term on the right can be rewritten as

$$\begin{aligned} e^{-\chi(x_0, \tilde{x}_\varepsilon^{(1)})} \int_0^{e^{\chi(x_0, \tilde{x}_\varepsilon^{(1)})}} F(\tilde{x}_\varepsilon^{(1)}, u) du \mathbf{1}_{(\chi(x_0, \tilde{x}_\varepsilon^{(1)}) \leq 0)} \\ + e^{-\chi(x_0, \tilde{x}_\varepsilon^{(1)})} \int_0^1 F(\tilde{x}_\varepsilon^{(1)}, u) du \mathbf{1}_{(\chi(x_0, \tilde{x}_\varepsilon^{(1)}) > 0)} \end{aligned}$$

while the second is

$$e^{-\chi(x_0, \tilde{x}_\varepsilon^{(1)})} \int_{e^{\chi(x_0, \tilde{x}_\varepsilon^{(1)})}}^1 F(\tilde{x}_\varepsilon^{(1)}, u) du \mathbf{1}_{(\chi(x_0, \tilde{x}_\varepsilon^{(1)}) \leq 0)}.$$

Combining terms we obtain

$$\begin{aligned} \mathbf{E} \left(\sum_{j=1}^{N_\varepsilon} F(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) \middle| x_0, \tilde{x}_\varepsilon^{(1)} \right) &= e^{-\chi(x_0, \tilde{x}_\varepsilon^{(1)})} \int_0^1 F(\tilde{x}_\varepsilon^{(1)}, u) du \mathbf{1}_{(\chi(x_0, \tilde{x}_\varepsilon^{(1)}) \leq 0)} \\ &\quad + e^{-\chi(x_0, \tilde{x}_\varepsilon^{(1)})} \int_0^1 F(\tilde{x}_\varepsilon^{(1)}, u) du \mathbf{1}_{(\chi(x_0, \tilde{x}_\varepsilon^{(1)}) > 0)} \\ &= e^{-\chi(x_0, \tilde{x}_\varepsilon^{(1)})} \int_0^1 F(\tilde{x}_\varepsilon^{(1)}, u) du, \end{aligned}$$

which establishes (4.2). A similar (but simpler) argument proves the result for \mathbf{E} replaced by \mathbf{E}^1 . \square

With Lemma 4.2 in hand we are ready to prove Theorem 4.1.

Proof of Theorem 4.1. First notice that by Lemma (4.2), we have the identity

$$\mathbf{E} \sum_{j=1}^{N_\varepsilon} F(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) = \mathbf{E} \left(e^{-\chi(x_0, y_\varepsilon)} \int_0^1 F(y_\varepsilon, u) du \right), \quad (4.3)$$

for any function $F(x, \theta)$. In particular, the required identity holds for $k = 1$ by setting $F(x, \theta) = f(x)$.

Now we proceed by induction, assuming that expression (4.1) holds for step $k - 1$ and prove that the relation also holds at k . Notice that, by the Markov property of the process,

$$\mathbf{E} \sum_{j=1}^{N_{k\varepsilon}} f(x_{k\varepsilon}^{(j)}) = \mathbf{E} \left(\sum_{j=1}^{N_\varepsilon} \mathbf{E}_{x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}} \left(\sum_{j=1}^{N_{(k-1)\varepsilon}} f(x_{(k-1)\varepsilon}^{(j)}) \right) \right).$$

Setting

$$F(x, \theta) = \mathbf{E}_{x, \theta} \sum_{j=1}^{N_{(k-1)\varepsilon}} f(x_{(k-1)\varepsilon}^{(j)})$$

in expression (4.3), we obtain

$$\begin{aligned} \mathbf{E} \sum_{j=1}^{N_{k\varepsilon}} f(x_{k\varepsilon}^{(j)}) &= \mathbf{E} \sum_{j=1}^{N_\varepsilon} F(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) \\ &= \mathbf{E} \left(e^{-\chi(x_0, y_\varepsilon)} \int_0^1 \mathbf{E}_{y_\varepsilon, u} \left(\sum_{j=1}^{N_{(k-1)\varepsilon}} f(x_{(k-1)\varepsilon}^{(j)}) \right) du \right). \end{aligned}$$

According to the rule for generating the initial ticket we have that

$$\int_0^1 \mathbf{E}_{y_\varepsilon, u} \left(\sum_{j=1}^{N_{(k-1)\varepsilon}} f(x_{(k-1)\varepsilon}^{(j)}) \right) du = \mathbf{E}_{y_\varepsilon} \sum_{j=1}^{N_{(k-1)\varepsilon}} f(x_{(k-1)\varepsilon}^{(j)})$$

and we have therefore shown that

$$\mathbf{E} \sum_{j=1}^{N_{k\varepsilon}} f(x_{k\varepsilon}^{(j)}) = \mathbf{E} \left(e^{-\chi(x_0, y_\varepsilon)} \mathbf{E}_{y_\varepsilon} \left(\sum_{j=1}^{N_{(k-1)\varepsilon}} f(x_{(k-1)\varepsilon}^{(j)}) \right) \right).$$

We can conclude, by our induction hypothesis, that

$$\begin{aligned} \mathbf{E} \sum_{j=1}^{N_{k\varepsilon}} f(x_{k\varepsilon}^{(j)}) &= \mathbf{E} \left(e^{-\chi(x_0, y_\varepsilon)} \mathbf{E}_{y_\varepsilon} \left(f(y_{(k-1)\varepsilon}) e^{-\sum_{k=0}^{t/\varepsilon-2} \chi(y_{k\varepsilon}, y_{(k+1)\varepsilon})} \right) \right) \\ &= \mathbf{E} \left(f(y_t) \exp \left(- \sum_{k=0}^{t/\varepsilon-1} \chi(y_{k\varepsilon}, y_{(k+1)\varepsilon}) \right) \right), \end{aligned}$$

and the proof is complete. A similar argument proves the result for \mathbf{E} replaced by \mathbf{E}^1 . \square

The second main result of this section compares the variance of the estimators generated by Algorithms 1 and 2. We have that

Theorem 4.3 *For all bounded functions f , one has the inequality*

$$\mathbf{var} \hat{f}_t \leq \mathbf{var}^1 \hat{f}_t.$$

The inequality is strict when f is strictly positive.

The following straightforward observation is an important tool in comparing Algorithms 1 and 2.

Lemma 4.4 *After replacing the rules*

$$\theta_{(k+1)\varepsilon}^{(j,1)} = \frac{\theta_{k\varepsilon}^{(j)}}{P^{(j)}} \quad \text{and} \quad \theta_{(k+1)\varepsilon}^{(j,i)} \sim \mathcal{U}((P^{(j)})^{-1}, 1) \quad \text{for } i > 1,$$

in Step 4 of Algorithm 2 by

$$\theta_{(k+1)\varepsilon}^{(j,i)} \sim \mathcal{U}(0, 1) \quad \text{for all } i, \tag{4.4}$$

the process generated by Algorithm 2 becomes identical in law to the process generated by Algorithm 1.

Remark 4.5 Note that the resampling of the tickets in (4.4) applies to *all* particles, including $i = 1$.

As with the proof of Theorem 4.1, the proof of Theorem 4.3 is inductive. The most cumbersome component of that argument is obtaining a comparison of cross terms that appear when expanding the variance of the estimators produced by Algorithms 1 and 2. This is encapsulated in the following lemma.

Lemma 4.6 *Let F be any non-negative bounded function of $\mathbb{R}^d \times \mathbb{R}$, which is decreasing in its second argument. Then*

$$\mathbf{E} \left(\sum_{j=1}^{N_\varepsilon} \sum_{i \neq j}^{N_\varepsilon} F(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) F(x_\varepsilon^{(i)}, \theta_\varepsilon^{(i)}) \right) \leq \mathbf{E}^1 \left(\sum_{j=1}^{N_\varepsilon} \sum_{i \neq j}^{N_\varepsilon} F(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) F(x_\varepsilon^{(i)}, \theta_\varepsilon^{(i)}) \right),$$

and the bound is strict when, for each x , $F(x, \theta)$ is a strictly decreasing function of θ .

Proof. First note that $x_\varepsilon^{(j)} = \tilde{x}_\varepsilon^{(1)}$ for each j . Furthermore, conditioned on x_0 , $\tilde{x}_\varepsilon^{(1)}$, and N_ε , the tickets are independent of each other and for $j \geq 2$ they are identically distributed (for Algorithm 1 they are identically distributed for $j \geq 1$). These facts imply that, for the new scheme,

$$\begin{aligned} \mathbf{E} \left(\sum_{j=1}^{N_\varepsilon} \sum_{i \neq j}^{N_\varepsilon} F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(j)}) F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(i)}) \middle| \tilde{x}_\varepsilon^{(1)} \right) = \\ \mathbf{E} \left(2 \mathbf{1}_{(N_\varepsilon > 1)} (N_\varepsilon - 1) \middle| \tilde{x}_\varepsilon^{(1)} \right) \mathbf{E} \left(F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(1)}) \middle| \tilde{x}_\varepsilon^{(1)} \right) \mathbf{E} \left(F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(2)}) \middle| \tilde{x}_\varepsilon^{(1)} \right) \\ + \mathbf{E} \left(\mathbf{1}_{(N_\varepsilon > 1)} (N_\varepsilon - 1)(N_\varepsilon - 2) \middle| \tilde{x}_\varepsilon^{(1)} \right) \mathbf{E} \left(F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(2)}) \middle| \tilde{x}_\varepsilon^{(1)} \right)^2. \end{aligned} \quad (4.5)$$

For Algorithm 1 the tickets $\theta_\varepsilon^{(j)}$ are i.i.d. conditioned on $\tilde{x}_\varepsilon^{(1)}$ so that

$$\begin{aligned} \mathbf{E}^1 \left(\sum_{j=1}^{N_\varepsilon} \sum_{i \neq j}^{N_\varepsilon} F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(j)}) F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(i)}) \middle| \tilde{x}_\varepsilon^{(1)} \right) = \\ \mathbf{E}^1 \left(N_\varepsilon (N_\varepsilon - 1) \middle| \tilde{x}_\varepsilon^{(1)} \right) \mathbf{E}^1 \left(F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(1)}) F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(1)}) \middle| \tilde{x}_\varepsilon^{(1)} \right)^2. \end{aligned} \quad (4.6)$$

Let

$$P = e^{-\chi(x_0, \tilde{x}_\varepsilon^{(1)})}.$$

Note that on the set $\{P \leq 1\}$, we have $N_\varepsilon \leq 1$, so that expressions (4.5) and (4.6) vanish. On the set $\{P > 1\}$, we have

$$\begin{aligned} \mathbf{E} \left(F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(2)}) \middle| \tilde{x}_\varepsilon^{(1)} \right) &= \frac{1}{1 - P^{-1}} \int_{P^{-1}}^1 F(\tilde{x}_\varepsilon^{(1)}, u) du, \\ \mathbf{E} \left(F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(1)}) \middle| \tilde{x}_\varepsilon^{(1)} \right) &= P \int_0^{P^{-1}} F(\tilde{x}_\varepsilon^{(1)}, u) du, \end{aligned}$$

and

$$\mathbf{E}^1 \left(F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(1)}) \middle| \tilde{x}_\varepsilon^{(1)} \right) = \int_0^1 F(\tilde{x}_\varepsilon^{(1)}, u) du,$$

so that

$$\mathbf{E}^1 \left(F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(1)}) \middle| \tilde{x}_\varepsilon^{(1)} \right) = \frac{1}{P} \mathbf{E} \left(F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(1)}) \middle| \tilde{x}_\varepsilon^{(1)} \right) + \frac{(P - 1)}{P} \mathbf{E} \left(F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(2)}) \middle| \tilde{x}_\varepsilon^{(1)} \right).$$

In other words, if we set

$$A = \mathbf{E}(F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(2)}) | \tilde{x}_\varepsilon^{(1)})$$

and

$$B = \mathbf{E}(F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(1)}) | \tilde{x}_\varepsilon^{(1)}) - \mathbf{E}(F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(2)}) | \tilde{x}_\varepsilon^{(1)}),$$

we have the identity

$$\mathbf{E}^1(F(\tilde{x}_\varepsilon^{(1)}, \theta_\varepsilon^{(1)}) | \tilde{x}_\varepsilon^{(1)}) = A + \frac{1}{P}B.$$

In Algorithm 2, $\theta_\varepsilon^{(1)} \leq \theta_\varepsilon^{(2)}$ so that, if F is a decreasing function of θ , then

$$B \geq 0.$$

Now note that the variable N_ε is the same for both algorithms and is given explicitly by the formula

$$N_\varepsilon = \mathbf{1}_{(\theta_\varepsilon^{(1)} \leq e^{-\chi(x_0, \tilde{x}_\varepsilon^{(1)})} \leq 1)} + \mathbf{1}_{(\chi(x_0, \tilde{x}_\varepsilon^{(1)}) < 0)} \\ \times \left(\left\lfloor e^{-\chi(x_0, \tilde{x}_\varepsilon^{(1)})} \right\rfloor + \mathbf{1}_{(U < e^{-\chi(x_0, \tilde{x}_\varepsilon^{(1)})} - \left\lfloor e^{-\chi(x_0, \tilde{x}_\varepsilon^{(1)})} \right\rfloor)} \right),$$

where U is a uniform random variable in $[0, 1]$. Using this formula we can compute the expectations involving N_ε explicitly. Defining

$$R = P - \lfloor P \rfloor,$$

on the event $\{P > 1\}$ we have that

$$\mathbf{E}^1(N_\varepsilon(N_\varepsilon - 1) | \tilde{x}_\varepsilon^{(1)}) = (P - R)(P - 1 + R), \\ \mathbf{E}(\mathbf{1}_{(N_\varepsilon > 1)}(N_\varepsilon - 1) | \tilde{x}_\varepsilon^{(1)}) = P - 1,$$

and

$$\mathbf{E}(\mathbf{1}_{(N_\varepsilon > 1)}(N_\varepsilon - 1)(N_\varepsilon - 2) | \tilde{x}_\varepsilon^{(1)}) = (P - 1 - R)(P - 2 + R).$$

The difference of terms (4.6) and (4.5), which vanishes on $P < 2$, can now be written as

$$2(P - 1)AB - 2(P - R)(P - 1 + R)AB \frac{1}{P} - (P - R)(P - 1 + R)B^2 \frac{1}{P^2}$$

on the set $\{P \geq 2\}$. Recalling that $B \geq 0$ (we can drop the last term) and rearranging we see that this expression is bounded above by

$$2 \frac{AB}{P} (P(P - 1) - (P - R)(P - 1 + R)). \quad (4.7)$$

Note that since $0 \leq R \leq 1$ we have $P - 1 \leq P - 1 + R \leq P - R \leq P$. The difference in the parenthesis in (4.7) is the difference in the area of two squares with equal perimeter, the second of which (in the sense of the last sequence of inequalities) closer to square. The difference is therefore non-positive. All bounds are strict whenever, for each x , $F(x, \theta)$ is a strictly decreasing function of θ . \square

Finally we complete the proof that the variance of the estimator generated by Algorithm 2 is bounded above by the variance of the estimator generated by Algorithm 1.

Proof of Theorem 4.3. By Theorem 4.1 we have that

$$\mathbf{E} \widehat{f}_t = \mathbf{E}^1 \widehat{f}_t ,$$

so it suffices to show that

$$\mathbf{E}(\widehat{f}_t)^2 \leq \mathbf{E}^1(\widehat{f}_t)^2 . \quad (4.8)$$

Furthermore, since the variance does not change by adding a constant to f we can, and will from now on, assume that $f(x) \geq 0$ for every x . In order to prove (4.8), we will proceed by induction. Since the random variables N_ε and $\tilde{x}_\varepsilon^{(1)}$ are the same in both schemes, the result is true for $k = 1$. Now assume that (4.8) holds through step $k - 1$. We will show that

$$\mathbf{E} \left(\sum_{j=1}^{N_{k\varepsilon}} f(x_{k\varepsilon}^{(j)}) \right)^2 \leq \mathbf{E}^1 \left(\sum_{j=1}^{N_{k\varepsilon}} f(x_{k\varepsilon}^{(j)}) \right)^2 . \quad (4.9)$$

Observe that we can write

$$\sum_{j=1}^{N_{k\varepsilon}} f(x_{k\varepsilon}^{(j)}) = \sum_{j=1}^{N_\varepsilon} \sum_{i=1}^{N_{\varepsilon,k\varepsilon}^{(j)}} f(x_{\varepsilon,k\varepsilon}^{(j,i)}) ,$$

where we used $N_{\varepsilon,k\varepsilon}^{(j)}$ to denote the number of particles at time $k\varepsilon$ whose ancestor at time ε is $x_\varepsilon^{(j)}$. The descendants of $x_\varepsilon^{(j)}$ are enumerated by $x_{\varepsilon,k\varepsilon}^{(j,i)}$ for $i = 1, 2, \dots, N_{\varepsilon,k\varepsilon}^{(j)}$.

By the conditional independence of the descendants of the particles at time $k\varepsilon$ when conditioned on x_0 , $\tilde{x}_\varepsilon^{(1)}$, and N_ε , we have that

$$\begin{aligned} \mathbf{E} \left(\sum_{j=1}^{N_{k\varepsilon}} f(x_{k\varepsilon}^{(j)}) \right)^2 &= \mathbf{E} \left(\sum_{j=1}^{N_\varepsilon} \mathbf{E}_{x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}} \left(\sum_{i=1}^{N_{(k-1)\varepsilon}} f(x_{(k-1)\varepsilon}^{(i)}) \right)^2 \right) \\ &\quad + \mathbf{E} \left(\sum_{j=1}^{N_\varepsilon} \sum_{i \neq j}^{N_\varepsilon} \mathbf{E}_{x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}} \left(\sum_{l=1}^{N_{(k-1)\varepsilon}} f(x_{(k-1)\varepsilon}^{(l)}) \right) \right. \\ &\quad \left. \times \mathbf{E}_{x_\varepsilon^{(i)}, \theta_\varepsilon^{(i)}} \left(\sum_{l=1}^{N_{(k-1)\varepsilon}} f(x_{(k-1)\varepsilon}^{(l)}) \right) \right) . \end{aligned} \quad (4.10)$$

An identical expression (with \mathbf{E} replaced by \mathbf{E}^1) holds for Algorithm 1. Setting

$$F_1(x, \theta) = \mathbf{E}_{x, \theta} \sum_{l=1}^{N_{(k-1)\varepsilon}} f(x_{(k-1)\varepsilon}^{(l)}) ,$$

and

$$F_2(x, \theta) = \mathbf{E}_{x, \theta} \left(\sum_{i=1}^{N_{(k-1)\varepsilon}} f(x_{(k-1)\varepsilon}^{(i)}) \right)^2 ,$$

we can rewrite the last identity in the case of Algorithm 2 as

$$\mathbf{E} \left(\sum_{j=1}^{N_{k\varepsilon}} f(x_{k\varepsilon}^{(j)}) \right)^2 = \mathbf{E} \sum_{j=1}^{N_\varepsilon} F_2(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) + \mathbf{E} \sum_{j=1}^{N_\varepsilon} \sum_{i \neq j}^{N_\varepsilon} F_1(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) F_1(x_\varepsilon^{(i)}, \theta_\varepsilon^{(i)}) .$$

In the case of Algorithm 1 with

$$F_1^0(x, \theta) = \mathbf{E}_{x, \theta}^1 \sum_{l=1}^{N_{(k-1)\varepsilon}} f(x_{(k-1)\varepsilon}^{(l)})$$

and

$$F_2^0(x, \theta) = \mathbf{E}_{x, \theta}^1 \left(\sum_{i=1}^{N_{(k-1)\varepsilon}} f(x_{(k-1)\varepsilon}^{(i)}) \right)^2,$$

we have

$$\mathbf{E}^1 \left(\sum_{j=1}^{N_{k\varepsilon}} f(x_{k\varepsilon}^{(j)}) \right)^2 = \mathbf{E}^1 \sum_{j=1}^{N_\varepsilon} F_2^0(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) + \mathbf{E}^1 \sum_{j=1}^{N_\varepsilon} \sum_{i \neq j}^{N_\varepsilon} F_1^0(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) F_1^0(x_\varepsilon^{(i)}, \theta_\varepsilon^{(i)}).$$

Note that Lemma 4.2 implies that

$$\mathbf{E} \sum_{j=1}^{N_\varepsilon} F_2(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) = \mathbf{E}^1 \sum_{j=1}^{N_\varepsilon} F_2(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}).$$

Under the rules of Algorithm 1, conditioned on x_0 and $\tilde{x}_\varepsilon^{(1)}$, the ticket $\theta_\varepsilon^{(j)}$ is independent of N_ε . So we can write the last equality as

$$\mathbf{E} \sum_{j=1}^{N_\varepsilon} F_2(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) = \mathbf{E}^1 \left(\sum_{j=1}^{N_\varepsilon} \mathbf{E}_{x_\varepsilon^{(j)}} \left(F_2(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) \right) \right).$$

By our inductive hypothesis we then have that

$$\mathbf{E} \sum_{j=1}^{N_\varepsilon} F_2(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) \leq \mathbf{E}^1 \left(\sum_{j=1}^{N_\varepsilon} \mathbf{E}_{x_\varepsilon^{(j)}} \left(F_2^0(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) \right) \right)$$

Appealing again to the conditional independence of $\theta_\varepsilon^{(j)}$ and N_ε under Algorithm 1 we have that

$$\mathbf{E} \sum_{j=1}^{N_\varepsilon} F_2(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) \leq \mathbf{E}^1 \sum_{j=1}^{N_\varepsilon} F_2^0(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}).$$

We now move on to the second term in (4.10). It follows from the definition of Algorithm 2 that the function F_1 is strictly decreasing in θ , so that Lemma 4.6 yields

$$\mathbf{E} \left(\sum_{j=1}^{N_\varepsilon} \sum_{i \neq j}^{N_\varepsilon} F_1(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) F_1(x_\varepsilon^{(i)}, \theta_\varepsilon^{(i)}) \right) \leq \mathbf{E}^1 \left(\sum_{j=1}^{N_\varepsilon} \sum_{i \neq j}^{N_\varepsilon} F_1(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) F_1(x_\varepsilon^{(i)}, \theta_\varepsilon^{(i)}) \right).$$

Under the rules of Algorithm 1, conditional on x_0 and $\tilde{x}_\varepsilon^{(1)}$, the $\theta_\varepsilon^{(1)}$ are all independent of each other and of N_ε . Therefore we can integrate over the tickets to obtain

$$\begin{aligned} & \mathbf{E} \left(\sum_{j=1}^{N_\varepsilon} \sum_{i \neq j}^{N_\varepsilon} F_1(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) F_1(x_\varepsilon^{(i)}, \theta_\varepsilon^{(i)}) \right) \\ & \leq \mathbf{E}^1 \left(\sum_{j=1}^{N_\varepsilon} \sum_{i \neq j}^{N_\varepsilon} \mathbf{E}_{x_\varepsilon^{(j)}} \left(F_1(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) \right) \mathbf{E}_{x_\varepsilon^{(i)}} \left(F_1(x_\varepsilon^{(i)}, \theta_\varepsilon^{(i)}) \right) \right). \end{aligned}$$

By Theorem 4.1 we have that

$$\mathbf{E}_{x_\varepsilon^{(j)}} \sum_{l=1}^{N_{(k-1)\varepsilon}} f(x_{(k-1)\varepsilon}^{(l)}) = \mathbf{E}_{x_\varepsilon^{(j)}}^0 \sum_{l=1}^{N_{(k-1)\varepsilon}} f(x_{(k-1)\varepsilon}^{(l)})$$

or

$$\mathbf{E}_{x_\varepsilon^{(j)}} F_1(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) = \mathbf{E}_{x_\varepsilon^{(j)}}^1 F_1^0(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}).$$

This yields

$$\begin{aligned} & \mathbf{E} \left(\sum_{j=1}^{N_\varepsilon} \sum_{i \neq j}^{N_\varepsilon} F_1(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) F_1(x_\varepsilon^{(i)}, \theta_\varepsilon^{(i)}) \right) \\ & \leq \mathbf{E}^1 \left(\sum_{j=1}^{N_\varepsilon} \sum_{i \neq j}^{N_\varepsilon} \mathbf{E}_{x_\varepsilon^{(j)}}^1 \left(F_1^0(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) \right) \mathbf{E}_{x_\varepsilon^{(i)}}^1 \left(F_1^0(x_\varepsilon^{(i)}, \theta_\varepsilon^{(i)}) \right) \right). \end{aligned}$$

Reinserting the tickets we obtain

$$\mathbf{E} \left(\sum_{j=1}^{N_\varepsilon} \sum_{i \neq j}^{N_\varepsilon} F_1(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) F_1(x_\varepsilon^{(i)}, \theta_\varepsilon^{(i)}) \right) \leq \mathbf{E}^1 \left(\sum_{j=1}^{N_\varepsilon} \sum_{i \neq j}^{N_\varepsilon} F_1^0(x_\varepsilon^{(j)}, \theta_\varepsilon^{(j)}) F_1^0(x_\varepsilon^{(i)}, \theta_\varepsilon^{(i)}) \right),$$

which completes the proof. \square

5 The continuous-time limit and the Brownian fan

From now on, we restrict ourselves to the analysis of the case $\chi(x, y) = V(y) - V(x)$ for some “potential” V defined on the state space of the underlying Markov process. We argue that if the underlying process is obtained by approximating a diffusion process then, unlike in the case of Algorithm 1, our modification, Algorithm 2, converges to a non-trivial limiting process as the stepsize ε converges to 0.

We will first provide a heuristic argument showing what kind of limiting process one would expect to obtain. The remainder of this article will then be devoted to rigorously constructing the limiting process and proving convergence in the simple case in which the underlying Markov chain is a random walk (rescaled so that it converges to a standard Brownian motion) and the potential V is linear. In this case the limiting process is a very natural object that does not seem to have been studied in the literature so far. We call this object, which is closely related to the construction in [Hut09], the *Brownian fan* (see Section 5.2). It also has a flavour very similar to the construction of the Brownian web [FINR04] and the Brownian net [SS08], although there does not seem to be an obvious transformation linking these objects.

5.1 Heuristic derivation of the continuous-time limit

Throughout this section the underlying Markov chain will be given by the following approximation to a diffusion:

$$y_{(k+1)\varepsilon} = y_{k\varepsilon} + \varepsilon F(y_{k\varepsilon}) + \sqrt{\varepsilon} \Sigma(y_{k\varepsilon}) \xi_{k+1}, \quad y_{k\varepsilon} \in \mathbf{R}^n, \quad (5.1)$$

where the ξ_k are a sequence of i.i.d. (not necessarily Gaussian) random variables with law ν and the identity on \mathbf{R}^n as their covariance matrix. The functions F and Σ are

sufficiently “nice” functions, but since this section is only heuristic, we do not state specific regularity, growth or non-degeneracy assumptions.

Concerning the function χ , we take $\chi(x, y) = V(y) - V(x)$ for some regular potential $V : \mathbf{R}^n \rightarrow \mathbf{R}$. Recall now that, as long as a particle is alive, its ticket θ evolves under Algorithm 2 as

$$\theta_{k\varepsilon}^{(j)} = \theta_{k\varepsilon}^{(j)} \exp(V(\tilde{x}_{(k+1)\varepsilon}^{(j)}) - V(x_{k\varepsilon}^{(j)})) .$$

It is natural therefore to replace θ by the quantity v given by

$$\exp(-v_t^{(j)}) = \theta_t^{(j)} \exp(-V(x_t^{(j)})) .$$

In this way, the new “tag” v does not change over time, but is assigned to a particle at the moment of its birth. Translating Steps 3 and 4 of the algorithm into this slightly different setting, we see that if a particle performs a step from x to y such that $V(y) < V(x)$, then it can potentially spawn one or more descendants. The tags v of the descendants are then distributed according to

$$e^{-v} \stackrel{\text{law}}{=} \mathcal{U}(e^{-V(x)}, e^{-V(y)}) , \quad (5.2)$$

and a particle with tag v lives as long as it stays within the region $\{x : V(x) \leq v\}$.

5.1.1 Description of the limit

For very small values of ε , the process described above has the following features. Taking the limit $\varepsilon \rightarrow 0$ in (5.1), we observe that each particle follows a diffusion process, solving the equation

$$dy_t = F(y_t) dt + \Sigma(y_t) dB_t , \quad (5.3)$$

where B_t is a standard d -dimensional Brownian motion. If the particle has tag v , then this process is killed as soon as it exits the sublevel set $\{x : V(x) \leq v\}$.

Consider the following representation of the object produced by Algorithm 2. Denote by $\mathbf{Q}_{x,v}^\varepsilon$ the law of the ε -discretization of (5.3) generated by (5.1) starting at x and killed upon exiting the set $\{y : V(y) \leq v\}$. Let τ and $\{w_{k\varepsilon}\}_{k=0}^{\tau/\varepsilon}$ be, respectively, the lifetime and trajectory of the original particle. The trajectories of the offspring of this initial trajectory are very nearly given (it will be true in the small ε limit) by a realisation $\mu^{\varepsilon,1}$ of a Poisson point process with intensity

$$\mathcal{Q}^\varepsilon(w, \cdot) = \varepsilon \sum_{k=0}^{\tau/\varepsilon-1} A^\varepsilon(k) \Theta_{k\varepsilon}^* \int \mathbf{Q}_{w_{k\varepsilon}, V(w_{k\varepsilon})+\delta}^\varepsilon \eta(k, d\delta) , \quad (5.4)$$

where, according to the rule for generating new offspring in Algorithm 2,

$$A^\varepsilon(k) = \begin{cases} e^{-(V(w_{(k+1)\varepsilon}) - V(w_{k\varepsilon}))} - 1 & \text{if } V(w_{(k+1)\varepsilon}) < V(w_{k\varepsilon}) \\ 0 & \text{if } V(w_{(k+1)\varepsilon}) \geq V(w_{k\varepsilon}) \end{cases}$$

and, according to the rule for generating offspring tickets in Algorithm 2,

$$\int f(\delta) \eta(k, d\delta) = \int_0^1 f(-\log(1 + u(e^{-(V(w_{(k+1)\varepsilon}) - V(w_{k\varepsilon}))} - 1))) du .$$

Here Θ_t is the map that shifts trajectories forward by time t . Since each offspring behaves independently just like the original particle, this suggests that the n th generation $\mu^{\varepsilon, n}$ of offspring is obtained recursively as a realisation of the Poisson point process with intensity given by

$$\mathcal{G}^{\varepsilon, n}(\cdot) = \int \mathcal{Q}^{\varepsilon}(\tilde{w}, \cdot) \mu^{\varepsilon, n-1}(d\tilde{w}) .$$

At each “microscopic” step, the probability of creating a descendent is of order $\sqrt{\varepsilon}$ so that, in the limit $\varepsilon \rightarrow 0$, each particle spawns descendants at infinite rate. However, any such descendant is created at distance $\mathcal{O}(\sqrt{\varepsilon})$ of the “barrier” $V(x) = v$. As a consequence, the probability that it survives for a time of order 1 before being killed is itself only of order $\sqrt{\varepsilon}$. Therefore, the rate at which a particle creates descendants that actually survive for a time τ of order 1 is finite, but tends to infinity as $\tau \rightarrow 0$.

Now we will consider the small ε -limit of the object we have constructed. The trajectory w_t becomes a sample path of (5.3) exiting the set $\{y : V(y) \leq v\}$ at time τ . Denote by $\mathbf{Q}_{x,v}$ the law of the diffusion (5.3) starting at x and killed upon exiting the set $\{y : V(y) \leq v\}$, which is a probability measure on some space of excursions in \mathbf{R}^n . The characterisation of the standard Itô excursion measure (see for example [RY91, Theorem 4.1] and [PY82]) then suggests that, for every $x \in \mathbf{R}^n$ such that $\nabla V \neq 0$ and Σ is non-degenerate, the limit

$$\mathbf{Q}_x = \lim_{\delta \rightarrow 0^+} \frac{1}{\delta} \mathbf{Q}_{x, V(x)+\delta} ,$$

exists as a σ -finite measure in the sense that $\frac{1}{\delta} \mathbf{Q}_{x, V(x)+\delta}$ restricted to the set of excursions longer than a fixed length converges weakly to \mathbf{Q}_x restricted to the same set.

The discussion so far suggests that for the limiting object, the trajectories of the first generation of offspring are given by a realisation μ^1 of the Poisson point process with intensity measure

$$\mathcal{Q}(w, \cdot) = \int_0^\tau A(w_t) \Theta_t^* \mathbf{Q}_{w_t} dt , \quad (5.5)$$

for some intensity $A : \mathbf{R}^n \rightarrow \mathbf{R}_+$ yet to be determined and that the n th generation μ^n of offspring is obtained recursively as a realisation of a Poisson point process with intensity given by

$$\mathcal{G}^n(\cdot) = \int \mathcal{Q}(\tilde{w}, \cdot) \mu^{n-1}(d\tilde{w}) ,$$

with \mathcal{Q} as in (5.5). In order to fully characterise the limiting object, it remains to provide an expression for the intensity function A .

Let us start by replacing $\mathbf{Q}_{x,v}^\varepsilon$ in equation (5.4) by $\mathbf{Q}_{x,v}$, i.e. by assuming that for small ε , excursions of the discrete process are very similar to excursions of its continuous time limit. We then apply the relations

$$\mathbf{Q}_{x, V(x)+\delta} \approx \delta \mathbf{Q}_x \quad (5.6)$$

and

$$V(w_{(k+1)\varepsilon}) - V(w_{k\varepsilon}) \approx \sqrt{\varepsilon} \langle \nabla V(w_{k\varepsilon}), \Sigma(w_{k\varepsilon}) \xi_{k+1} \rangle$$

with the ξ_{k+1} as in (5.1). We then formally obtain

$$\begin{aligned} \mathcal{Q}^\varepsilon(w, \cdot) &\approx \varepsilon \sum_{k=0}^{\tau/\varepsilon-1} \mathbf{1}_{\langle \nabla V(x), \Sigma(x)\xi_{k+1} \rangle < 0} \langle \nabla V(w_{k\varepsilon}), \Sigma(w_{k\varepsilon})\xi_{k+1} \rangle \\ &\quad \times \int_0^1 u \langle \nabla V(w_{k\varepsilon}), \Sigma(w_{k\varepsilon})\xi_{k+1} \rangle du \Theta_{k\varepsilon}^* \mathbf{Q}_{w_{k\varepsilon}}. \end{aligned}$$

Our arguments so far therefore suggest that

$$A(x) = \frac{1}{2} \int_{\langle \nabla V(x), \Sigma(x)z \rangle < 0} \langle \nabla V(x), \Sigma(x)z \rangle^2 \nu(dz), \quad (5.7)$$

where the distribution ν has mean 0 and identity covariance matrix. Assuming that ν is symmetric this becomes

$$A(x) = c \langle \nabla V(x), \Sigma(x) \Sigma^T(x) \nabla V(x) \rangle, \quad (5.8)$$

where $c = \frac{1}{4}$. If ν is not symmetric, one might even expect a prefactor c that depends on x .

In fact, as we will see in a specific case in the remainder of this section, the correct value is $c = \frac{1}{2}$, whether the law of ξ is symmetric or not. The reason for this discrepancy is that the relation

$$\mathbf{Q}_{x, V(x)+\delta}^\varepsilon \approx \delta \mathbf{Q}_x$$

used in our derivation is only valid if $\delta \gg \sqrt{\varepsilon}$. In our case however, one precisely has $\delta \sim \sqrt{\varepsilon}$, which introduces a correction factor that eventually gives rise to the value $c = \frac{1}{2}$. The aim of the next subsection is to show in more detail how this factor $\frac{1}{2}$ arises in the simplest situation where $F = 0$ and $\Sigma = 1$.

5.1.2 The case of Brownian motion

We now consider the one-dimensional case, where the limiting underlying process is simple Brownian motion. Regarding the underlying discrete problem, we consider the Markov chain defined recursively by

$$y_{(k+1)\varepsilon} = y_{k\varepsilon} + \sqrt{\varepsilon} \xi_{k+1}, \quad (5.9)$$

for an i.i.d. sequence of centred random variables ξ_k with law ν and variance 1. For the potential function V , we choose $V(x) = -ax$ for some $a > 0$.

In order to show that the constant c appearing in (5.8) is equal to $\frac{1}{2}$, we will now argue that if we denote by \mathbf{Q} the standard Itô excursion measure (which we normalise in such a way that $\mathbf{Q} = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \mathbf{Q}_\varepsilon$, where \mathbf{Q}_ε is the law of a standard Brownian motion starting at ε and killed when it hits the origin) and by \mathbf{Q}_z^ε the law of the random walk (5.9) starting at $\sqrt{\varepsilon}z$ and stopped as soon as it takes negative values, then there exists a function G such that

$$\mathbf{Q}_z^\varepsilon \approx \sqrt{\varepsilon} G(z) \mathbf{Q}, \quad (5.10)$$

as $\varepsilon \rightarrow 0$ when both sides are restricted to excursions that survive for at least some fixed amount of time. We will see that the function G behaves like $G(z) \approx z$ for large values of z , but has a non-trivial behaviour for values of order 1. In terms of our notation from

the previous subsection (since y_t is spatially homogeneous and V is linear) this implies that the approximation (5.6) should really have been replaced by

$$\mathbf{Q}_{x, V(x)+\delta}^\varepsilon \approx a\sqrt{\varepsilon} G\left(\frac{\delta}{a\sqrt{\varepsilon}}\right) \mathbf{Q}.$$

Since we assumed $a > 0$, our process creates offspring only when it performs a step towards the right, i.e. when $\xi_{k+1} > 0$. The probability that a new particle is created in the k -th step is approximately $a\xi_{k+1}$. Furthermore, the small ε rule for the generation of tags implied by Algorithm 2 is

$$\frac{\delta}{a\sqrt{\varepsilon}} = \frac{a^{-1}v + x}{\sqrt{\varepsilon}} \sim \mathcal{U}(0, \xi_{k+1}). \quad (5.11)$$

As a consequence, once we have identified the function G in (5.10), our arguments in the previous section lead to the formula

$$A(x) = \int_0^\infty (az) \left(\frac{1}{z} \int_0^z aG(y) dy \right) \nu(dz)$$

where ν is the law of the steps ξ_k . If we can show that

$$\int_0^\infty \int_0^z G(y) dy \nu(dz) = \frac{1}{2}, \quad (5.12)$$

then we will have

$$A(x) = \frac{a^2}{2},$$

a formula consistent with a choice of $c = \frac{1}{2}$ in (5.8).

In order to identify G , we note that if a random walk starting from $\sqrt{\varepsilon}\delta$ survives for some time of order 1 before becoming negative then, with overwhelming probability, it will have reached a height of at least $\varepsilon^{1/4}$ (say). Furthermore, if we condition the random walk \mathbf{Q}_z^ε to reach a level $\sqrt{\varepsilon}\gamma$ with $1 \ll \gamma \ll \varepsilon^{-1/2}$, one would expect its law to be well approximated by $\sqrt{\varepsilon}\gamma \mathbf{Q}$ when restricted to excursions that survive for a time of order 1.

As a consequence, we expect that

$$\mathbf{Q}_z^\varepsilon \approx \bar{P}_{z,\gamma} \sqrt{\varepsilon} \gamma \mathbf{Q}, \quad \gamma \gg 1,$$

where $\bar{P}_{z,\gamma}$ denotes the probability that the simple random walk (5.9) with $\varepsilon = 1$ started at z reaches the level γ before becoming negative.

The remainder of this section is devoted to the proof of the fact that if we define $\bar{P}_{z,\gamma}$ in this way, then under some integrability assumptions for the one-step probability ν , the limit

$$G(z) = \lim_{\gamma \rightarrow \infty} \gamma \bar{P}_{z,\gamma},$$

exists and does indeed satisfy (5.12), independently of the choice of ν . Actually, we will prove these statements for the quantity $P_{z,\gamma} = \bar{P}_{z,\gamma+z}$, which we interpret as the probability that the random walk starting at the origin reaches $[\gamma, \infty)$ before reaching $(-\infty, -z]$. Our first result is as follows:

Proposition 5.1 *Assume that the law ν satisfies $\nu(\{|x| \geq K\}) \leq C \exp(-cK^\beta)$ for all $K \geq 0$ and some strictly positive constants c , C and β . Then, the limit*

$$G(s) = \lim_{\gamma \rightarrow \infty} \gamma P_{s,\gamma} \quad (5.13)$$

exists and satisfies the relations

$$G(s) = \int_{-s}^{\infty} G(s+z) \nu(dz), \quad s \geq 0, \quad \lim_{s \rightarrow \infty} \frac{G(s)}{s} = 1. \quad (5.14)$$

Furthermore, for every $\delta > 0$ there exists C such that the bound

$$|(\gamma + s)P_{s,\gamma} - G(s)| \leq C \frac{1+s}{\gamma^{\frac{1}{2}-\delta}}, \quad (5.15)$$

holds uniformly for all $s \geq 0$ and $\gamma \geq 1 \vee s$.

Proof. Denote by y_k the k th step of the random walk starting at the origin. Our main tool is the quantitative convergence result [Fra73], which states that the supremum distance between a Wiener process and the diffusively rescaled random walk over n steps is of order $n^{-1/4}$.

As a consequence we claim first that, for every $\delta > 0$ there exists a constant C such that, for every $a \in [\frac{1}{3}, 3]$, we have the bound

$$\left| P_{a\gamma,\gamma} - \frac{a}{1+a} \right| \leq \frac{C}{\gamma^{\frac{1}{2}-\delta}}, \quad (5.16)$$

valid for every $\gamma \geq 1$. Indeed, for any $n \geq 1$, it follows from the previously quoted convergence result that there exists a Brownian motion B such that $|B_t - y_{\lfloor t \rfloor}| \leq n^{1/4+\delta}$ for all $t \in [0, n]$ with probability greater than $1 - C/n^q$. Here, $\delta > 0$ and $q \geq 1$ are arbitrary, but the constant C of course depends on them.

Take n such that $n^{1/4+\delta} \leq \gamma$. If y hits $[\gamma, \infty)$ before $(-\infty, -a\gamma]$, then either $\sup_{t \leq n} |B_t - y_{\lfloor t \rfloor}| > n^{1/4+\delta}$, or $\sup_{t \leq n} |B_t| \leq 3\gamma$, or B hits $[\gamma - n^{1/4+\delta}, \infty)$ before it hits $(-\infty, -a\gamma - n^{1/4+\delta}]$. As a consequence,

$$P_{a\gamma,\gamma} \leq \frac{a\gamma + n^{1/4+\delta}}{(1+a)\gamma} + \frac{C}{n^q} + \exp(-cn/\gamma^2).$$

Reversing the roles of γ and $a\gamma$, we thus obtain the bound

$$\left| P_{a\gamma,\gamma} - \frac{a}{1+a} \right| \lesssim \frac{n^{1/4+\delta}}{\gamma} + \frac{1}{n^q} + \exp(-cn/\gamma^2).$$

Choosing δ small enough and $n = \gamma^{2+\delta}$, the claim (5.16) then follows.

In order to obtain the convergence of the right hand side in (5.13), we make use of the fact that, for $\bar{\gamma} > \gamma$, one has the identity

$$P_{s,\bar{\gamma}} = P_{s,\gamma} \int_0^\infty P_{s+\gamma+z,\bar{\gamma}-\gamma-z} \nu_\gamma(dz), \quad (5.17)$$

where ν_γ is the law of the “overshoot” $y_n - \gamma$ at the first time n such that $y_n \geq \gamma$, conditioned on never reaching below the level $-s$. Since $P_{s+\gamma+z,\bar{\gamma}-\gamma-z}$ is an increasing function of z , we immediately obtain the lower bound

$$P_{s,\bar{\gamma}} \geq P_{s,\gamma} P_{s+\gamma,\bar{\gamma}-\gamma},$$

If we choose $\gamma = a\bar{\gamma}$ for $a \in [\frac{1}{4}, \frac{1}{2}]$, it then follows from (5.16) that

$$P_{s,\bar{\gamma}} \geq P_{s,\gamma} \left(\frac{\gamma + s}{\bar{\gamma} + s} - \frac{C}{\gamma^{\frac{1}{2}-\delta}} \right),$$

for all γ sufficiently large and uniformly over all $s \in [0, \gamma]$. Setting $Q_{s,\gamma} = (\gamma + s)P_{s,\gamma}$, it thus follows that one has the bound

$$Q_{s,\bar{\gamma}} \geq Q_{s,\gamma} \left(1 - \frac{C}{\gamma^{\frac{1}{2}-\delta}} \right), \quad (5.18)$$

possibly for a different constant C . Let $\gamma_0 \geq 1$ be such that the factor on the right of this equation is greater than $1/2$. By (5.16), there then exists $s_0 \geq \gamma_0$ such that for $s \geq s_0$ and $\gamma \in [s, 2s]$ one has $Q_{s,\gamma} \geq C(1 + s)$. Furthermore, for $s \leq s_0$ and $\gamma \in [(1 \vee s), 2(\gamma_0 \vee s)]$, there exists a non-zero constant such that $Q_{s,\gamma} \geq C$. Iterating (5.18), we then conclude that there exists a constant $C > 0$ such that the bound

$$Q_{s,\gamma} \geq C(1 + s),$$

holds uniformly over all $s > 0$ and all $\gamma \geq 1 \vee s$.

On the other hand, for arbitrary $\alpha > 0$, one has from (5.17) the lower bound

$$P_{s,\bar{\gamma}} \leq (\nu_\gamma(\{x > \gamma^\alpha\}) + P_{s+\gamma+\gamma^\alpha, \bar{\gamma}-\gamma-\gamma^\alpha})P_{s,\gamma}.$$

In order to bound $\nu_\gamma(\{x > \gamma^\alpha\})$, we note that this event can happen only if either one of the first γ^3 increments exceeds γ^α , or the random walk never exceeds the value γ within these γ^3 steps. Similarly to before, it then follows that

$$\nu_\gamma(\{x > \gamma^\alpha\}) \lesssim \frac{1}{P_{s,\gamma}} (\gamma^3 \exp(-c\gamma^{\alpha\beta}) + \gamma^{-q} + \exp(-c\gamma)),$$

for every $q > 0$ and uniformly over $s \leq \gamma$. It follows from the lower bound on $P_{s,\gamma}$ obtained previously that $\nu_\gamma(\{x > \gamma^\alpha\}) \lesssim \gamma^{-q}$ for any power $q > 0$, so that we obtain the upper bound

$$P_{s,\bar{\gamma}} \leq P_{s,\gamma} \left(\frac{\gamma + s}{\bar{\gamma} + s} + \frac{C}{\gamma^{\frac{1}{2}-\delta}} \right), \quad (5.19)$$

with the same domain of validity as before. Using a very similar argument as before, we obtain a constant \bar{C} such that $\gamma P_{s,\gamma} \leq \bar{C}(1 + s)$ uniformly over $s > 0$ and $\gamma \geq 1 \vee s$.

Combining the bounds we just obtained, we obtain

$$|\bar{\gamma}P_{s,\bar{\gamma}} - \gamma P_{s,\gamma}| \leq \frac{C}{\gamma^{\frac{1}{2}+\delta}},$$

uniformly over $\bar{\gamma} > \gamma > s$, from which it follows immediately that the sequence $\{\gamma P_{s,\gamma}\}_{\gamma \geq 1}$ is Cauchy, so that it has a limit $G(s)$.

It remains to show that G has the desired properties. The first one follows immediately from the identity

$$\gamma P_{s,\gamma} = \int_{\mathbf{R}} \gamma P_{s+z, \gamma-z} \nu(dz),$$

which holds provided that we define the integrand to be 1 for $z > \gamma$ and 0 for $z < -s$.

In order to show that $G(s)/s \rightarrow 1$, we fix some (large) value s and choose $\gamma_n = 2^n s$. It then follows from (5.16) that

$$|Q_0 - s| \lesssim s^{\frac{1}{2}+\delta},$$

where we used the notation $Q_n = (\gamma_n + s)P_{s, \gamma_n}$ as a shorthand. Furthermore, it follows immediately from (5.19) that there exists a constant C independent of s such that $|Q_n| \leq Cs$ uniformly in n . As a consequence, we obtain the recursive bound

$$|Q_n - Q_{n-1}| \leq \frac{Cs}{\gamma_n^{\frac{1}{2}-\delta}}.$$

Summing over n yields $|Q_n - s| \lesssim s^{\frac{1}{2}+\delta}$, uniformly in n , so that the claim follows. The quantitative error bound (5.15) follows in the same way. \square

Corollary 5.2 *In the same setting as above, one has the bound*

$$\left| P_{s, \gamma} - \frac{s}{\gamma + s} \right| \lesssim \frac{1 + s^{\frac{1}{2}+\delta}}{\gamma},$$

uniformly for all $s \geq 0$ and $\gamma \geq 1 \vee s$.

Proof. Combine (5.15) with the bounds on $G(s) - s$ obtained at the end of the proof above. \square

Somewhat surprising is the fact that the function G obtained in the Proposition 5.1 does indeed satisfy (5.12), independently of the choice of transition probability ν , provided that we assume that ν has some exponential moment.

Proposition 5.3 *Let G be as in Proposition 5.1 and assume that the law ν satisfies*

$$\int_{\mathbf{R}} e^{c|z|} \nu(dz) < \infty,$$

for some $c > 0$. Then, one has the identity $\int_0^\infty \nu([s, \infty)) G(s) ds = \frac{1}{2}$.

Remark 5.4 Note that, by Fubini's theorem,

$$\int_0^\infty \nu([s, \infty)) G(s) ds = \int_0^\infty \int_0^s G(y) dy \nu(ds),$$

so that we do obtain (5.12).

Proof. Integrating (5.14) from 0 to an arbitrary value $K > 0$ and applying Fubini's theorem, we obtain the identity

$$\int_0^K G(s) ds = \int_0^\infty G(z) \nu([z - K, z]) dz. \quad (5.20)$$

In this proof we denote by $\mathcal{I} = \int_0^\infty G(z) \nu([z, \infty)) dz$ the quantity of interest. Simple algebraic manipulations then yield from (5.20)

$$\begin{aligned} \mathcal{I} &= \int_0^\infty G(z) \nu([z - K, \infty)) dz - \int_0^\infty G(z) \nu([z - K, z]) dz \\ &= \int_0^\infty G(z) \nu([z - K, \infty)) dz - \int_0^K G(z) dz \\ &= \int_0^\infty G(z) (\nu([z - K, \infty)) - \mathbf{1}_{z < K}) dz. \end{aligned}$$

Since this identity holds for every $K > 0$, it follows in particular that one has

$$\begin{aligned} \mathcal{I} &= \varepsilon \int_0^\infty G(z) \int_0^\infty (\nu([z - K, \infty)) - \mathbf{1}_{z < K}) e^{-\varepsilon K} dK dz \\ &= \int_0^\infty G(z) \left(\varepsilon \int_0^\infty \nu([z - K, \infty)) e^{-\varepsilon K} dK - e^{-\varepsilon z} \right) dz, \end{aligned} \quad (5.21)$$

for every $\varepsilon > 0$. At this stage, we note that one has the identity

$$\varepsilon \int_0^\infty \nu([z - K, \infty)) e^{-\varepsilon K} dK = e^{-\varepsilon z} \mathbf{E} e^{\varepsilon \xi} - \varepsilon e^{-\varepsilon z} \int_z^\infty e^{\varepsilon K} \nu([K, \infty)) dK,$$

where ξ denotes an arbitrary random variable with law ν . Since ν has some exponential moment by assumption, $\nu([K, \infty))$ decays exponentially so that the second term in this identity satisfies

$$\left| \varepsilon e^{-\varepsilon z} \int_z^\infty e^{\varepsilon K} \nu([K, \infty)) dK \right| \leq C \varepsilon e^{-\gamma z},$$

for some constants $\gamma, C > 0$, provided that ε is small enough. Inserting this into (5.21), it follows that

$$\mathcal{I} = \int_0^\infty G(z) e^{-\varepsilon z} \mathbf{E}(e^{\varepsilon \xi} - 1) dz + \mathcal{O}(\varepsilon).$$

At this stage, we use again the fact that ξ has exponential moments to deduce that

$$\mathbf{E} e^{\varepsilon \xi} - 1 = \frac{\varepsilon^2}{2} + \mathcal{O}(\varepsilon^3),$$

where we used the fact that $\mathbf{E} \xi = 0$ and $\mathbf{E} \xi^2 = 1$, so that

$$\mathcal{I} = \frac{\varepsilon^2}{2} \int_0^\infty G(z) e^{-\varepsilon z} dz + \mathcal{O}(\varepsilon).$$

It then follows from the fact that $\lim_{s \rightarrow \infty} G(s)/s = 1$ and the dominated convergence theorem that

$$\mathcal{I} = \lim_{\varepsilon \rightarrow 0} \frac{\varepsilon^2}{2} \int_0^\infty z e^{-\varepsilon z} dt = \frac{1}{2},$$

which is precisely the desired expression. \square

Remark 5.5 It is clear that these results should hold under much weaker integrability conditions on ν . However, since we need some exponential moments on ν at several places in the sequel, we did not try to improve on this.

5.2 Some properties of the limiting process

In this section, we provide a rigorous definition of the limiting process loosely defined in Section 5.1.1, and we study some of its properties. In order to be able to use existing results on Brownian excursions, we restrict ourselves to the same situation as in Section 5.1.2, namely the case where the underlying diffusion is a Brownian motion and the potential $V(x) = -ax$ is linear. We call the resulting object the *Brownian fan*.

5.2.1 Recursive Poisson point processes

Before we give a formal definition of the Brownian fan, we define a “recursive Poisson point process”. Loosely speaking, this is a Crump-Mode-Jagers process [Jag75] with Poisson distributed offspring, but where the number of offspring of any given individual is allowed to be almost surely infinite. Note again that our construction is very similar to the one given in [Hut09]. Given a Polish space \mathcal{X} and a function $F: \mathcal{X} \rightarrow \mathbf{R}_+$, we denote throughout this section by $M_+^F(\mathcal{X})$ the space of σ -finite measures μ on \mathcal{X} such that

$$\mu(F^{-1}(0)) = 0, \quad \mu(\{x : F(x) > \varepsilon\}) < \infty,$$

for all $\varepsilon > 0$. We endow this with the topology of convergence in total variation on each set of the form $\{x : F(x) > 1/n\}$. Given a (measurable) map \mathcal{Q} from \mathcal{X} to $M_+^F(\mathcal{X})$, we can then build for every $x \in \mathcal{X}$ a point process as follows.

Define $\mu_x^0 = \delta_x$ and, for $n \geq 1$, define μ_x^n recursively as a (conditionally independent of the μ_x^ℓ with $\ell < n$) realisation of a Poisson point process with intensity measure

$$\mathcal{Q}_n = \int_{\mathcal{X}} \mathcal{Q}(y) \mu_x^{n-1}(dy),$$

where we view μ_x^n as a random σ -finite positive integer-valued measure on \mathcal{X} . (In principle, it may happen that $\mathcal{Q}_n(\{x : F(x) > \varepsilon\}) = \infty$ for some $\varepsilon > 0$. In this case, our construction stops there.)

When then set

$$\mu_x^{[n]} = \sum_{\ell=0}^n \mu_x^\ell, \quad (5.22)$$

and we call $\mu_x^{[n]}$ the *recursive Poisson point process* of depth n with kernel \mathcal{Q} . We will occasionally need to refer the Brownian fan spawned by an initial Brownian motion w . For this purpose we will use the symbol $\mu_w^{[n]}$ (or μ_w^n for a specific generation) and rely on the context to differentiate $\mu_x^{[n]}$ and $\mu_w^{[n]}$. If, in these symbols we omit the subscript entirely then it is assumed that $x = 0$. In general, there is no reason to expect the sequence $\mu_x^{[n]}$ to converge to a finite limit. However, one has the following simple criterion ensuring that this is the case:

Lemma 5.6 *Let F and \mathcal{Q} be as above and assume that there exists $c < 1$ such that $\int F(y)\mathcal{Q}(x, dy) \leq cF(x)$ for every $x \in \mathcal{X}$. Then, for every $x \in \mathcal{X}$, there exists a random σ -finite measure $\mu_x^{[\infty]}$ on \mathcal{X} such that $\lim_{n \rightarrow \infty} \mathbf{E} \int F(y)(\mu_x^{[n]} - \mu_x^{[\infty]})(dy) = 0$.*

Proof. Fix $x \in \mathcal{X}$ and denote by \mathcal{F}_n the σ -algebra generated by $\mu_x^{[n]}$. It then follows from the definition of the $\mu_x^{[n]}$ that

$$\mathbf{E} \left(\int_{\mathcal{X}} F(y) \mu_x^{\ell+1}(dy) \middle| \mathcal{F}_\ell \right) = \int_{\mathcal{X}} \int_{\mathcal{X}} F(y) \mathcal{Q}(z, dy) \mu_x^\ell(dz) \leq c \int_{\mathcal{X}} F(z) \mu_x^\ell(dz).$$

As a consequence, one has $\mathbf{E} \int_{\mathcal{X}} F(y) \mu_x^\ell(dy) \leq c^\ell F(x)$, and the claim follows. \square

Remark 5.7 A useful identity is the following. Denote by $\{\tilde{\mu}_y^{[\infty]}\}_{y \in \mathcal{X}}$ a collection of *independent* copies of recursive Poisson point processes with “initial conditions” y . Then one has the identity in law

$$\mu_x^{[\infty]} \stackrel{\text{law}}{=} \delta_x + \int_{\mathcal{X}} \tilde{\mu}_y^{[\infty]} \mu_x^1(dy), \quad (5.23)$$

where, as before, μ_x^1 is a realisation of a Poisson point process with intensity $\mathcal{Q}(x, \cdot)$, which is itself independent of the $\tilde{\mu}_y^{[\infty]}$. This identity makes sense since the integral on the right is really just a countable sum.

5.2.2 Construction of the Brownian fan

We now denote by \mathcal{E} the set of excursions with values in \mathbf{R} . We consider elements of \mathcal{E} as triples (s, t, y) where $s < t \in \mathbf{R} \cup \{+\infty\}$, and $y \in \mathcal{C}(\mathbf{R}, \mathbf{R})$ has the property that $y_\tau = y_t$ for $\tau \geq t$ and $y_\tau = y_s$ for $\tau \leq s$. We also write \mathcal{E}_0 for the subset of those triples (s, t, y) such that $s = 0$.

Denoting a generic excursion by w , we write $s(w)$ for its starting time and $e(w)$ for its end time, *i.e.*, $s(s, t, y) = s$, and $e(s, t, y) = t$. We also denote by $l(w)$ the lifetime of the excursion, which is the interval $l(w) = [s(w), e(w)]$. In order to keep notations compact, we will also identify an excursion with its path component, making the abuse of notation $w_t = y_t$. There is a natural metric on \mathcal{E} given by

$$d(\tilde{w}, w) = d_l(\tilde{w}, w) \vee \sum_{k \geq 1} 2^{-k} \left(1 \wedge \sup_{|t| \leq 2^k} |\tilde{w}_t - w_t| \right), \quad (5.24)$$

where the distance d_l between the supports is given by

$$d_l(\tilde{w}, w) = 1 \wedge (|s(\tilde{w}) - s(w)| + |\tanh e(\tilde{w}) - \tanh e(w)|).$$

The reason for this particular choice of metric is that it ensures that \mathcal{E} is a Polish space, while still allowing for infinite excursions.

For $\tau \in \mathbf{R}$ and $v \in \mathcal{E}$, we denote by $\Theta_{\tau, v}: \mathcal{E} \rightarrow \mathcal{E}$ the shift map given by

$$\Theta_{\tau, v}: (s, t, w) \mapsto (s + \tau, t + \tau, w_{\cdot + \tau} + v_\tau),$$

which essentially changes the coordinate system so that the origin $(0, 0)$ is mapped to (τ, v_τ) . Denoting as before by \mathbf{Q} the standard Itô excursion measure, we now give the following definition:

Definition 5.8 The *Brownian fan* with intensity $a > 0$ is the recursive Poisson point process on \mathcal{E} with kernel

$$\mathcal{Q}(w, \cdot) = \frac{a}{2} \int_{s(w)}^{e(w)} \Theta_{\tau, w}^* \mathbf{Q} d\tau, \quad (5.25)$$

and initial condition given by a realisation of Brownian motion, starting at the origin and killed when it reaches the level $-L$, where L is exponentially distributed with mean a .

Remark 5.9 The reason for killing the original Brownian motion at this particular level is natural, due to the distribution of the initial tag in Step 1 of the algorithm. It is however essentially irrelevant to the mathematical construction.

Remark 5.10 Formally, the Brownian fan is a particular case of the Virgin Island Model [Hut09] with a playing the same role in both models, $h = a$, and $g = 1/2$. The differences are twofold. First, the case of constant non-vanishing a actually doesn't fall within the framework of [Hut09] since the author there uses the standing assumption




Figure 6: A typical realisation of the Brownian fan with intensity $a = \frac{3}{2}$ over the time interval $[0, 1]$. Successive generations are drawn in lighter shades of grey.

that $a(0) = 0$. The other difference is mostly one of perspective. While we have so far defined the Brownian fan as a point process on a space of excursions, one of the purposes of this article is to show that it is also well-behaved as an actual Markov process with values in a suitable state space of (possibly infinite) point configurations.

Remark 5.11 By only keeping track of the genealogy of the particles and not their precise locations, one can construct a “real tree” on top of which the Brownian fan is constructed (loosely speaking) by attaching a Brownian excursion to each branch. This is very similar in spirit to Le Gall’s construction of the Brownian map [LG07, LG10], starting from Aldous’s continuous random tree [Ald91]. The scaling properties of the Brownian fan however are quite different. In particular, Theorem 5.13 below implies that the underlying tree has Hausdorff dimension 1, as opposed to the CRT which has Hausdorff dimension 2.

Before we proceed, let us show that it is possible to verify the assumptions of Lemma 5.6, so that this object actually exists for every $a > 0$:

Proposition 5.12 *The kernel \mathcal{Q} defined in (5.25) satisfies the assumption of Lemma 5.6 with the choice*

$$F(w) = e^{-\eta s(w)}(1 - e^{-\eta |l(w)|}),$$

provided that η is large enough.

Proof. It follows from the properties of \mathbf{Q} that there exists a constant C independent of γ and a such that

$$\int_{\mathcal{E}} F(\tilde{w}) \mathcal{Q}(w, d\tilde{w}) = \frac{a}{\sqrt{2\pi}} \int_{s(\tilde{w})}^{\epsilon(\tilde{w})} e^{-\eta \tau} \int_0^\infty (1 - e^{-\eta s}) s^{-3/2} ds d\tau$$

$$\leq \frac{Ca}{\sqrt{\eta}}(e^{-\eta s(\tilde{w})} - e^{-\eta \mathfrak{e}(\tilde{w})}) = \frac{Ca}{\sqrt{\eta}}F(v),$$

where C is a constant independent of a and η . The claim then follows by choosing $\sqrt{\eta} > Ca$. \square

The remainder of this section is devoted to a study of the basic properties of the Brownian fan. In particular, we will show that there exists a suitable space \mathcal{X} of particle configurations such that it can be viewed as a \mathcal{X} -valued Markov process with continuous sample paths that satisfies the Feller property.

5.2.3 Number of particles and workload rate

Define the set $\mathcal{N}_t \subset \mathcal{E}$ of excursions that are “alive at time t ” by

$$\mathcal{N}_t = \{w \in \mathcal{E} : \mathfrak{s}(w) < t < \mathfrak{e}(w)\}.$$

With this notation, the number of particles alive at time t for the Brownian fan is given by

$$N_t = \mu^{[\infty]}(\mathcal{N}_t),$$

which is in principle allowed to be infinite.

Theorem 5.13 *There exist a constant $C > 0$ and a strictly positive continuous decreasing function $\lambda: \mathbf{R}_+ \rightarrow \mathbf{R}_+$ such that*

$$\mathbf{E} \exp(\lambda_t N_t) \leq C, \quad (5.26)$$

holds uniformly over all $t > 0$.

Remark 5.14 We will see in the proof that one can choose λ of the form

$$\lambda_t = K^{-1} e^{-Kt}, \quad (5.27)$$

for K sufficiently large.

Proof. For $\lambda > 0$ and $s, t \geq 0$, set

$$N_{s,t}^\lambda = \log \mathbf{E} \exp(\lambda \mu_w^{[\infty]}(\mathcal{N}_t)),$$

where w is any excursion starting from 0 with lifetime s . Since, by the definition (5.25), the value of $N_{s,t}^\lambda$ does not depend on the precise choice of excursion, we do not include it in the notation. It also follows from the construction of $\mu_w^{[\infty]}$ that the function $t \mapsto N_{s,t}^\lambda$ is increasing in t and that $N_{s,t}^\lambda = N_{t,t}^\lambda$ for $s \geq t$.

We also define M_t^λ by

$$M_t^\lambda = \log \mathbf{E} \exp\left(\lambda \int_{\mathcal{E}} \mu_w^{[\infty]}(\mathcal{N}_t) \mathbf{M}(dw)\right), \quad (5.28)$$

where \mathbf{M} is a Poisson random measure with intensity measure \mathbf{Q} and the realisations $\mu_w^{[\infty]}$ are independent of \mathbf{Q} and of each other. While N^λ measures the total number of offspring alive at time t due to an excursion starting at time 0, M^λ measures the rate at which these offspring are created.

Indeed, combining (5.23) with the definition of μ_w^1 and the superposition principle for Poisson point processes, we have the identity

$$N_{s,t}^\lambda = \lambda \mathbf{1}_{s \geq t} + a \int_0^{s \wedge t} M_{t-r}^\lambda dr. \quad (5.29)$$

It therefore remains to obtain suitable bounds on M_t^λ .

It follows from (5.29) and standard properties of Poisson point processes (see for example [PZ07, Theorem 6.3]) that one has the identity

$$\begin{aligned} M_t^\lambda &= \int_{\mathcal{E}} (\mathbf{E} \exp(\lambda \mu_w^{[\infty]}(\mathcal{N}_t)) - 1) \mathbf{Q}(dw) = \int_{\mathcal{E}} (e^{N_\lambda(\mathfrak{e}(w), t)} - 1) \mathbf{Q}(dw) \\ &\leq \int_{\mathcal{E}} \left(\exp\left(\lambda \mathbf{1}_{\mathcal{N}_t}(w) + a \int_0^{\mathfrak{e}(w) \wedge t} M_{t-s}^\lambda ds\right) - 1 \right) \mathbf{Q}(dw). \end{aligned}$$

At this stage, we note that the integrand appearing in this expression depends on w only through $\mathfrak{e}(w)$. It is then convenient to break the integral into a contribution coming from $\mathfrak{e}(w) > t$, as well as its complement. Since, under \mathbf{Q} , the quantity $\mathfrak{e}(w)$ is distributed according to the measure $\frac{s^{-3/2}}{\sqrt{2\pi}} ds$, this yields

$$\begin{aligned} M_t^\lambda &\leq \left(\exp\left(\lambda + a \int_0^t M_r^\lambda dr\right) - 1 \right) \mathbf{Q}(\mathfrak{e}(w) \geq t) \\ &\quad + \int_0^t \left(\exp\left(a \int_0^s M_{t-r}^\lambda dr\right) - 1 \right) \frac{s^{-3/2}}{\sqrt{2\pi}} ds. \end{aligned}$$

We now assume that both λ and t are sufficiently small so that

$$\lambda + a \int_0^t M_r^\lambda dr \leq 1. \quad (5.30)$$

This assumption allows us to use the bound $e^t - 1 \leq 2t$, so that we obtain the more manageable expression

$$\begin{aligned} M_t^\lambda &\leq \frac{4t^{-1/2}}{\sqrt{2\pi}} \left(\lambda + a \int_0^t M_r^\lambda dr \right) + 2a \int_0^t \int_0^s M_{t-r}^\lambda dr \frac{s^{-3/2}}{\sqrt{2\pi}} ds \\ &= \frac{4\lambda t^{-1/2}}{\sqrt{2\pi}} + \frac{4at^{-1/2}}{\sqrt{2\pi}} \int_0^t M_r^\lambda dr + 2a \int_0^t \int_s^t M_r^\lambda dr \frac{(t-s)^{-3/2}}{\sqrt{2\pi}} ds \\ &= \frac{4\lambda t^{-1/2}}{\sqrt{2\pi}} + \frac{4at^{-1/2}}{\sqrt{2\pi}} \int_0^t M_r^\lambda dr + \frac{4a}{\sqrt{2\pi}} \int_0^t (t-s)^{-1/2} M_s^\lambda ds \\ &\leq \frac{4\lambda t^{-1/2}}{\sqrt{2\pi}} + \frac{8a}{\sqrt{2\pi}} \int_0^t (t-s)^{-1/2} M_s^\lambda ds. \end{aligned}$$

Writing $H_t^\lambda = t^{1/2} M_t^\lambda$, we thus obtain the bound

$$H_t^\lambda \leq \frac{4\lambda}{\sqrt{2\pi}} + \frac{8a}{\sqrt{2\pi}} t^{1/2} \int_0^t (t-s)^{-1/2} s^{-1/2} H_s^\lambda ds. \quad (5.31)$$

We can now apply the fractional version of Gronwall's lemma [NR02, Lemma 7.6] (with $b = \frac{8a}{\sqrt{2\pi}}$, $a = \frac{4\lambda}{\sqrt{2\pi}}$, and $\alpha = \frac{1}{2}$), so that there exists a constant $C > 0$ depending on a but independent of λ such that

$$H_t^\lambda \leq C\lambda \exp(Ct).$$

From this, we immediately deduce from (5.29) and the definition of H^λ a similar bound on $N_{t,t}^\lambda$. Choosing $\lambda = K^{-1}e^{-Kt}$ for sufficiently large t then allows to satisfy (5.30) and to obtain $N_{t,t}^\lambda \leq 2$, thus completing the proof. \square

As a corollary, we obtain a rather sharp bound on the modulus of continuity of the total workload process $\mathcal{W}_t = \int_0^t N_s ds$. One has

Proposition 5.15 *For every $T > 0$, one has*

$$\sup_{t \leq T} \lim_{h \rightarrow 0} \frac{|\mathcal{W}_{t+h} - \mathcal{W}_t|}{h |\log h|} < \infty, \quad (5.32)$$

almost surely.

Proof. It follows from the generalised Young inequality that, for every $a, b \in \mathbf{R}_+$, and every $\lambda, \eta > 0$, one has the inequality

$$ab \leq \frac{\eta}{\lambda} (e^{\lambda a} - 1 - \lambda a + (1 + b/\eta) \log(1 + b/\eta) - b/\eta),$$

so that

$$N_t \leq \frac{\eta}{\lambda} (e^{\lambda N_t} + (1 + 1/\eta) \log(1 + 1/\eta)).$$

It follows immediately that

$$|\mathcal{W}_{t+h} - \mathcal{W}_t| \leq \frac{\eta}{\lambda} \int_t^{t+h} e^{\lambda N_s} ds + \frac{h(1+\eta)}{\lambda} \log(1 + 1/\eta).$$

Setting $\eta = h$, we obtain the bound

$$|\mathcal{W}_{t+h} - \mathcal{W}_t| \leq \frac{h}{\lambda} \int_0^{T+1} e^{\lambda N_s} ds + C_\lambda h |\log h|,$$

uniformly over all $h \leq 1$ and all $t \in [0, T]$. The claim now follows immediately from Theorem 5.13. \square

Although the number of particles alive at any deterministic time has exponential moments, there exists a *dense* set of exceptional times for which $N_t = \infty$. For one, this follows from the fact that, under \mathbf{Q} , $\mathfrak{e}(w)$ is distributed proportionally to $s^{-3/2} ds$, so that every particle creates an infinite number of offspring in every time interval.

Actually, one has the even stronger statement that there is a dense set of exceptional times at which the number of particles belonging to the *first* generation of offspring is infinite. Indeed, if we denote by \mathbf{M} a Poisson random measure on \mathbf{R}_+^2 with density $cs^{-3/2} dr ds$ for a suitable constant c , then the number N_t^1 of particles in the first generation of offspring is given by

$$N_t^1 = \mathbf{M}(A_t), \quad A_t = \{(r, s) \in [0, t] \times \mathbf{R}_+ : s \geq t - r\}.$$

For $k \geq 0$ and $d \in \{1, \dots, 2^k\}$, we then set

$$A_{k,d} = [(d-1)2^{-k}, d2^{-k}] \times [4^{-k}, 4^{1-k}],$$

so that, by the scaling properties of M , the random variables $N_{k,d} = M(A_{k,d})$ form a sequence of i.i.d. Poisson random variables. For any given point (r, s) , we set $B_{(r,s)} = [r, r + s]$, which is the set of times t such that $(r, s) \in A_t$, and we set

$$D_{(r,s)} = \{(k, d) : B_{(r',s')} \subset B_{(r,s)} \forall (r', s') \in A_{k,d}\} .$$

Since the set $D_{(r,s)}$ is infinite for every (r, s) , we can then build a sequence (r_n, s_n) recursively in the following way. Start with $(r_0, s_0) = (0, 1)$ and then, given (r_n, s_n) for some $n \geq 0$, define (k_n, d_n) as the first (in lexicographic order) element $(k, d) \in P_{(r_n, s_n)}$ such that $N_{k,d} \geq 1$. We then set (r_{n+1}, s_{n+1}) to be one of the points of M located in A_{k_n, d_n} . By construction, one then has $\cap_{n \geq 1} B_{(r_n, s_n)} = \{t\}$ for some $t \in [0, 1]$, and $M(A_t) \geq \sum_n M(A_{k_n, d_n}) = \infty$, as stated. Of course, the interval $[0, 1]$ in this procedure is arbitrary. If we want to show that there exists an exceptional time within any deterministic time interval $[t_0, t_1]$, it suffices to start the algorithm we just described with $r_0 = t_0$ and $s_0 = t_1 - t_0$.

5.3 The Brownian fan as a Markov process

In this section, we slightly shift our perspective. We do no longer consider our process as a point process of excursions, but we consider it as an evolving system of particles. Our system will therefore be described by a Markov process in some space of integer-valued measures on a subset of \mathbf{R}^2 corresponding to the admissible combinations of “position + tag”. The problem is that, as we have seen in the previous section, there are exceptional times at which the limiting process consists of infinitely many particles. The first challenge is therefore to construct a space \mathcal{X} of integer-valued measures with a sensible topology which can still accommodate these “bad” configurations in such a way that the limiting process is continuous both as a function of time and as a function of its initial configuration.

Once this space is defined, we show that the Brownian fan possesses the Feller property in \mathcal{X} (i.e. the corresponding Markov semigroup leaves the space of bounded continuous functions invariant). In fact, we will show that it preserves the space of Lipschitz continuous functions. The continuity property of the discrete-time process established below in Proposition 6.8 also implies the time continuity of the Brownian fan (see Theorem 6.11 below). These properties allow us to conclude that the Brownian fan $t \mapsto \mu_t$ is in fact a strong Markov process.

In Section 5.3.4 below, we will furthermore compute its generator A on a class of “nice” test functions. We will also present a “negative” result showing that if we denote by T_ε the one-step Markov operator corresponding to the evolution of Algorithm 2, then one has $A \neq \lim_{\varepsilon \rightarrow 0} \varepsilon^{-1}(T_\varepsilon - 1)$. This is in stark contrast with, for example, Euler approximations to stochastic differential equations, where such an equality would hold, at least when applied to sufficiently regular test functions.

5.3.1 State space

Our construction is essentially the Wasserstein-1 analogue of the construction given in [FG10]. Let $\mathcal{M} \subset \mathbf{R}^n$ be a convex open set with boundary $\partial\mathcal{M}$. For $p \in (0, 1]$, we then denote by $\ell^p(\mathcal{M})$ the set of all integer-valued measures μ on \mathcal{M} such that

$$\|\mu\|_p = \int_{\mathcal{M}} d^p(y, \partial\mathcal{M}) \mu(dy) < \infty , \quad (5.33)$$

where $d(y, \partial\mathcal{M})$ denotes the (Euclidean) distance from y to the boundary of \mathcal{M} . Note that since this quantity vanishes at the boundary, there are elements $\mu \in \ell^p(\mathcal{M})$ such that $\mu(\mathcal{M}) = \infty$.

We endow $\ell^p(\mathcal{M})$ with a slight modification of the Wasserstein-1 metric by setting:

$$\|\mu - \nu\|_p = \sup_{f \in \text{Lip}_p^0(\mathcal{M})} \left(\int f(y) \mu(dy) - \int f(y) \nu(dy) \right), \quad (5.34)$$

where we denoted by $\text{Lip}_p^0(\mathcal{M})$ the set of all functions $f: \mathcal{M} \rightarrow \mathbf{R}$ such that

$$|f(x) - f(y)| \leq |x - y|^p, \quad (5.35)$$

for all $x, y \in \mathcal{M}$, and $f(y) = 0$ for all $y \in \partial\mathcal{M}$.

Remark 5.16 Our notation is consistent in the sense that if we take for ν the null measure in (5.34), then we precisely recover (5.33). This can be seen by taking $f(x) = d^p(x, \partial\mathcal{M})$, which is optimal by (5.35) and the triangle inequality.

If μ and ν happen to have the same (finite) mass, then the expression (5.34) does not change when one adds a constant to f . In this case, we are thus reduced to the usual Wasserstein-1 distance between μ and ν , but with respect to the modified distance function

$$d_p(x, y) = |x - y|^p \wedge (d^p(x, \partial\mathcal{M}) + d^p(y, \partial\mathcal{M})).$$

Note that the completion of \mathcal{M} under the distance function d_p consists of $\mathcal{M} \cup \{\Delta\}$, where Δ is a single “point on the boundary” such that $d_p(x, \Delta) = d^p(x, \partial\mathcal{M})$ for every $x \in \mathcal{M}$.

If one has $\mu(\mathcal{M}) < \nu(\mathcal{M}) < \infty$, then the distance $\|\cdot\|_p$ reduces to the Wasserstein-1 distance (again with respect to d_p) between $\bar{\mu}$ and ν , where $\bar{\mu}$ is obtained from μ by placing a mass $\nu(\mathcal{M}) - \mu(\mathcal{M})$ on the boundary Δ . The following alternative characterisation of (5.34) in the case of purely atomic measures is a version of the Monge-Kantorovich duality in this context:

Lemma 5.17 *Consider a situation where $\mu = \sum_{i=1}^N \delta_{x_i}$ and $\nu = \sum_{i=1}^M \delta_{y_i}$. Then,*

$$\|\mu - \nu\|_p = \inf_{\sigma \in S_{N+M}} \sum_{i=1}^{N+M} d^p(x_i, y_{\sigma(i)}),$$

where S_{N+M} is the group of permutations of $N + M$ elements and we set $x_j = \Delta$ for $j > N$ and $y_j = \Delta$ for $j > M$.

Proof. See for example [Pra08]. □

This characterisation suggests the following “interpolation” procedure between elements in $\ell^p(\mathcal{M})$. Let $\mu = \sum_{i=1}^N \delta_{x_i}$ and $\nu = \sum_{i=1}^N \delta_{y_i}$, where we assumed that both measures charge the same number of points (this is something that we can always achieve by possibly adding points on Δ). assume furthermore that these points are ordered in such a way that

$$\|\mu - \nu\|_p = \sum_{i=1}^N |x_i - y_i|^p.$$

Again, this can always be enforced by suitably reordering the points and possibly adding points on the boundary. We then define, for $t \in (0, 1)$, the “linear interpolation” $L_t(\mu, \nu)$ by

$$L_t(\mu, \nu) = \sum_{i=1}^N \delta_{z_i}, \quad z_i = ty_i + (1-t)x_i. \quad (5.36)$$

Note that this procedure is not necessarily unique, but it is easy to resolve this ambiguity by optimising over the possible pairings $\{(x_i, y_i)\}$ realising the above construction, according to some arbitrary criteria.

In any case, one can check that this construction then has the property that

$$\|L_s(\mu, \nu) - L_t(\mu, \nu)\|_p \leq |t - s|^p \|\mu - \nu\|_p, \quad (5.37)$$

for any $s, t \in [0, 1]$, which will be a useful fact in the sequel.

5.3.2 Definition of the process

For the remainder of this section, we set

$$\mathcal{M} = \{(x, v) \in \mathbf{R}^2 : v > -ax\},$$

which is the natural configuration space for our process. We will use capital letters to distinguish points in \mathcal{M} from points in \mathbf{R} . By Theorem 5.13, we already know that, for any fixed time t , the Brownian fan almost surely has only finitely many particles alive at time t . Define now the evaluation map $E_t: \mathcal{E} \rightarrow \mathcal{M} \cup \{\Delta\}$ by

$$E_t(w) = \begin{cases} (w_t, -aw_{\mathfrak{s}(w)}) & \text{if } t \in \mathfrak{l}(w), \\ \Delta & \text{otherwise.} \end{cases}$$

For a given “initial condition” $(x, v) \in \mathcal{M}$, we then set

$$\mu_t = E_t^* \mu_w^{[\infty]},$$

which is an $\ell^p(\mathcal{M})$ -valued random variable. Here, w is a realisation of a Brownian motion starting at x and killed when it hits $-v/a$, and $\mu_w^{[\infty]}$ is the corresponding realisation of the Brownian fan. (Just so that E_t has the correct effect on w , one can for example set $\mathfrak{s}(w) = -1$ and make sure that $w(-1) = -v/a$.) As a consequence of Theorem 5.13 and of our definition of the Brownian fan, we then indeed have $\mu_t \in \ell^p(\mathcal{M})$ for every $p \leq 1$.

Note at this stage that we can simply discard the, typically infinite, mass on Δ by identifying measures that only differ on Δ . As already mentioned earlier, this is consistent with the identification $\Delta \sim \partial\mathcal{M}$ already made in the interpretation of the construction of $\ell^p(\mathcal{M})$.

This construction can be extended to any initial condition in $\ell^p(\mathcal{M})$ with finite total mass, by considering independent Brownian fans for each particle. As a consequence of the Markov property of the Brownian excursion and the independence properties of Poisson point processes, it is then straightforward to verify that $t \mapsto \mu_t$ is indeed a Markov process.

Actually, by Proposition 6.2, we know that for any fixed collection of deterministic times $\{t_1, \dots, t_k\}$, one has $\mu_{t_k}(\mathcal{M}) < \infty$ almost surely, so that our construction determines a probability measure on $(\ell^p(\mathcal{M}))^{\mathbf{R}^+}$ by Kolmogorov’s extension theorem. At this stage however, we know absolutely nothing about the continuity properties of this process, and this is the subject of the remainder of this subsection.

5.3.3 Feller property

We now show that the Brownian fan constructed in Sections 5.2.2 and 5.3.2 has the Feller property in $\ell^p(\mathcal{M})$ for every $p \leq 1$.

As in [FG10], we could have defined spaces $\ell^p(\mathcal{M})$ in a natural way for $p > 1$. However, the Feller property would fail in this case because of the following simple heuristic argument. For any p , we can change the initial condition by an amount less than δ in ℓ^p by creating N particles at distance $\varepsilon = (\delta/N)^{1/p}$ from the boundary of \mathcal{M} . For ε small, the probability that any such particle survives up to time 1 (say) is bounded from below by $c\varepsilon$ for some $c > 0$. On average, the number of survivors will thus be of the order of $\varepsilon N \sim \delta^{1/p} N^{(p-1)/p}$. Furthermore, at time $1/2$, each of these surviving particles will be at a distance of order 1 of the boundary of \mathcal{M} . As a consequence, by increasing N but keeping δ fixed (or even sending δ to 0 sufficiently slowly), the law of the process at time 1 with an initial condition arbitrarily close to 0 can be at arbitrarily large distance of 0, so that the Feller property fails.

For $p \leq 1$ on the other hand, we have

Proposition 5.18 *For any $p \leq 1$, the Brownian fan gives rise to a Feller process in $\ell^p(\mathcal{M})$. Even more, the corresponding Markov semigroup preserves the space of bounded Lipschitz continuous functions.*

Proof. For any two initial conditions μ_0 and $\bar{\mu}_0$, write as before

$$\mu_0 = \sum_{j=1}^N \delta_{X_0^{(j)}} , \quad \bar{\mu}_0 = \sum_{j=1}^N \delta_{\bar{X}_0^{(j)}} ,$$

with the $X_0^{(j)} \in \mathcal{M}$ and $\bar{X}_0^{(j)} \in \mathcal{M}$ chosen in such a way that

$$\|\mu_0 - \bar{\mu}_0\|_p = \sum_{j=1}^N d_p(X_0^{(j)}, \bar{X}_0^{(j)}) . \quad (5.38)$$

Our aim now is to show that there exists a constant C such that

$$\mathbf{E} \|\mu_t - \bar{\mu}_t\|_p \leq C \|\mu_0 - \bar{\mu}_0\|_p ,$$

independently of $t \leq 1$, where the pair $(\mu_t, \bar{\mu}_t)$ is any coupling between the Brownian fans starting from μ_0 and $\bar{\mu}_0$ respectively. Denote by $\mu_t^{(j)}$ the contribution to μ_t originating from the initial particle $X_0^{(j)}$ and similarly for $\bar{\mu}_t^{(j)}$. Then, by the triangle inequality, one obtains the bound

$$\mathbf{E} \|\mu_t - \bar{\mu}_t\|_p \leq \sum_{j \geq 1} \mathbf{E} \|\mu_t^{(j)} - \bar{\mu}_t^{(j)}\|_p ,$$

so that the claim follows from (5.38) if we can show that

$$\mathbf{E} \|\mu_t^{(j)} - \bar{\mu}_t^{(j)}\|_p \leq C d_p(X_0^{(j)}, \bar{X}_0^{(j)}) .$$

In other words, it suffices to consider the special case when both μ_0 and $\bar{\mu}_0$ consist of one single particle, which we denote by $X_0 = (x_0, v_0)$ and $\bar{X}_0 = (\bar{x}_0, \bar{v}_0)$ respectively.

One then constructs a coupling between the two processes μ_t and $\bar{\mu}_t$ by running both particles with the same Brownian motion and spawning children according to the

same Poisson process (as long as the corresponding particle is alive). We denote by X_t and \bar{X}_t the evolutions of the two initial particles in \mathcal{M} , driven by the same realisation of a Brownian motion, and stopped when they reach $\partial\mathcal{M}$. We can assume without loss of generality that $v_0 + ax_0 < \bar{v}_0 + a\bar{x}_0$, so that the particle X dies before the particle \bar{X} . Denoting by τ and $\bar{\tau}$ the respective lifetimes of these particles, one thus has $\tau \leq \bar{\tau}$.

Denote now by M_t the (random) measure on $[0, t] \times \mathcal{M}$ which is such that, for $I \subset [0, t]$ and $A \subset \mathcal{M}$, $M_t(I \times A)$ is the number of particles in A at time t that are offspring of a particle created from the “ancestor particle” X or \bar{X} at some time $s \in I$. With this notation, if we denote by $\Xi: \mathcal{M} \rightarrow \mathcal{M}$ the map

$$\Xi(x, v) = (x + x_0 - \bar{x}_0, v - a(x_0 - \bar{x}_0)) ,$$

then one has the decompositions

$$\begin{aligned} \mu_t(A) &= \mathbf{1}_{\tau \geq t} \delta_{X_t}(A) + \Xi^* M_t([0, \tau \wedge t] \times A) , \\ \bar{\mu}_t(A) &= \mathbf{1}_{\bar{\tau} \geq t} \delta_{\bar{X}_t}(A) + M_t([0, \bar{\tau} \wedge t] \times A) . \end{aligned}$$

Denote now by δ the *Euclidean* distance between the two initial particles, so that their ℓ^p -distance is δ^p . It then follows immediately from the above decomposition that one has the bound

$$\begin{aligned} \|\mu_t - \bar{\mu}_t\|_p &\leq \delta^p + \mathbf{1}_{t \in [\tau, \bar{\tau}]} d_p(\bar{X}_t, \partial\mathcal{M}) + \delta^p M_t([0, \tau \wedge t] \times \mathcal{M}) \\ &\quad + \int_{\mathcal{M}} d_p(Y, \partial\mathcal{M}) M_t([\tau \wedge t, \bar{\tau} \wedge t] \times dY) . \end{aligned}$$

Since $d(\bar{X}(\tau), \partial\mathcal{M}) = \delta$ by the definition of τ and δ , it follows from Jensen’s inequality and the Martingale property of (stopped) Brownian motion that one has the bound $\mathbf{E} \mathbf{1}_{t \in [\tau, \bar{\tau}]} d_p(\bar{X}_t, \partial\mathcal{M}) \leq \delta^p$. It also follows from Theorem 5.13 that $\mathbf{E} M_t([0, \tau \wedge t] \times \mathcal{M}) < \infty$, independently of δ . Finally, it follows from an argument very similar to the proof of Theorem 5.13 that

$$\mathbf{E} \int_{\mathcal{M}} d_p(Y, \partial\mathcal{M}) M_t(ds \times dY) \leq C ds ,$$

uniformly over $s \in [0, t]$. It follows that

$$\mathbf{E} \int_{\mathcal{M}} d_p(Y, \partial\mathcal{M}) M_t([\tau \wedge t, \bar{\tau} \wedge t] \times dY) \leq C \mathbf{E} |\bar{\tau} \wedge t - \tau \wedge t| \leq C \delta ,$$

where we used the fact that if τ_δ is the first hitting time of 0 by a Brownian motion starting at δ , then $\mathbf{E}(\tau_\delta \wedge 1) \leq C\delta$. Combining these bounds completes the proof. \square

5.3.4 Lack of convergence of the generators

One standard method to prove convergence of a sequence of Markov processes to a limiting process, once tightness has been established, is to show that the corresponding generators converge in a suitable sense. In our situation, one actually does *not* expect the generator of the approximate process to converge to that of the limiting process, when testing it on “nice” test functions. We first argue at a technical level why this is the case, before providing an intuitive explanation.

Inspired by [EK86, Daw93], we consider test functions of the form

$$F(\mu) = \exp(\langle \log f, \mu \rangle) , \tag{5.39}$$

where $f: \mathcal{M} \rightarrow \mathbf{R}_+$ is a sufficiently smooth function such that $f(x, v) = 1$ for $(x, v) \in \partial\mathcal{M}$. This boundary condition is required since elements $\mu \in \ell^p(\mathcal{M})$ can have infinite mass (and, as we have already seen, the limiting process really does acquire infinite mass at some exceptional times) accumulating near $\partial\mathcal{M}$. Being in $\ell^p(\mathcal{M})$ for $p \leq 1$ does however ensure that smooth functions such as above are integrable.

Exploiting the independence structure of the process as well as its space homogeneity, we can reduce ourselves to the case of an initial condition of the form $\mu_0 = \delta_{(x,v)}$ for some $v < -ax$. In this case, for sufficiently small $\varepsilon > 0$, the probability that the original particle dies within the time interval ε is of the order ε^p for any p . We therefore only need to take into account the possibility of creating some descendant(s), with the killing mechanism being taken care of by the boundary condition of f . While the average number of “second generation” descendants is of order ε , any such descendant will typically have travelled to a distance of order $\sqrt{\varepsilon}$ from $\partial\mathcal{M}$, so that only the first generation has a chance of contributing to the generator.

We then have

$$\frac{1}{\varepsilon}(\mathbf{E}F(\mu_\varepsilon) - F(\mu_0)) \approx A_0 f(x, v) + \frac{f(x, v)}{\varepsilon} \mathbf{E}\langle f - 1, M_\varepsilon \rangle ,$$

where

$$A_0 = \frac{1}{2} \partial_x^2 ,$$

is the generator of Brownian motion, and where M_ε is the (projection to time ε of the) Poisson point process yielding the first generation of offsprings. Note now that since these offspring will be created near $\partial\mathcal{M}$ and since $f = 1$ there, we can further approximate this expression by

$$\frac{1}{\varepsilon}(\mathbf{E}F(\mu_\varepsilon) - F(\mu_0)) \approx A_0 f(x, v) + \frac{f(x, v)}{\varepsilon} f'(x) \mathbf{E} \int_{\mathcal{M}} (\bar{x} + a^{-1}\bar{v}) M_\varepsilon(d\bar{x}, d\bar{v}) .$$

Here, we wrote $f'(x)$ as a shortcut for $\partial_x f(x, v)|_{v=-ax}$. Denoting by e_s the position, relative from its starting point, of an excursion of length s and making use of the formula (5.25) for the intensity measure of M_ε , we obtain for the last term in this equation the expression

$$\frac{1}{\varepsilon} \mathbf{E} \int_{\mathcal{M}} (\bar{x} + a^{-1}\bar{v}) M_\varepsilon(d\bar{x}, d\bar{v}) = \frac{a}{2\sqrt{2\pi\varepsilon}} \int_0^\infty \int_0^{s \wedge \varepsilon} \mathbf{E} e_s(t) dt s^{-3/2} ds . \quad (5.40)$$

At this stage we note that, for a Brownian excursion of length s , we have for $t \leq s$ the identity

$$\mathbf{E} e_s(t) = \sqrt{\frac{8t}{\pi s}} (s - t) .$$

Inserting this into (5.40), a tedious but straightforward calculation then yields

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \mathbf{E} \int_{\mathcal{M}} (\bar{x} + a^{-1}\bar{v}) M_\varepsilon(d\bar{x}, d\bar{v}) = \frac{a}{2} ,$$

so that we finally obtain for the generator A the expression

$$AF(\mu_0) = A_0 f(x, v) + \frac{a}{2} f(x, v) f'(x) . \quad (5.41)$$

Recall that this is for the particular case where $\mu_0 = \delta_{(x,v)}$. In the general case, we can use the independence structure of the process to obtain

$$AF(\mu_0) = F(\mu_0) \int_{\mathcal{M}} \left(\frac{A_0 f(x, v)}{f(x, v)} + \frac{a}{2} f'(x) \right) \mu_0(dx, dv) .$$

Remark 5.19 Compare this with the generator of a usual branching diffusion, where the term $f'(x)$ would be replaced by $a(f(x) - 1)$, with a the branching rate.

On the other hand, if we denote by T_ε the Markov operator describing one step of Algorithm 2, we might expect that one also obtains A as the limit $\varepsilon^{-1}(T_\varepsilon - 1)$ as $\varepsilon \rightarrow 0$. This would indeed be the case if there was no branching or if branching did only occur at a finite rate. Considering again initial conditions of the form $\mu_0 = \delta_{(x,v)}$ for some $v < -ax$. By reasoning similar to that in the beginning of Section 5, we obtain

$$\varepsilon^{-1}(T_\varepsilon F - F)(\mu_0) \approx A_0 f(x, v) + \frac{a}{2} f(x, v) f'(x) \int_0^\infty y^2 \nu(dy), \quad (5.42)$$

which is always *different* from (5.41), and is actually what we would have obtained from the wrong guess (5.7).

A possible reason for this discrepancy is that, while the Markov semigroup of the Brownian fan does indeed preserve test functions of the type (5.39), we do not expect this to be true of the Markov operator T_ε . Instead, the “correct” space of test functions for T_ε is of the same type, but the function f should have a “boundary layer” near $\partial\mathcal{M}$.

Remark 5.20 Another reason why the generator of the Brownian fan is not such a useful object is that many seemingly innocent observables, like for example the total number \mathcal{N} of particles, do *not* belong to its domain. This follows from the fact that if we consider again a simple initial condition μ_0 as above, then $\mathbf{E}\mathcal{N}(\mu_\varepsilon) - 1 \approx \sqrt{\varepsilon}$ for small ε . The reason why the total number of particles nevertheless remains finite (at least for fixed times) is that there are exceptional states where one or more particles are very near $\partial\mathcal{M}$ and for which $\mathbf{E}\mathcal{N}(\mu_\varepsilon) - 1 \approx -\mathcal{O}(1)$.

6 Convergence to the Brownian fan

The aim of this section is to provide a rigorous proof of the fact that, in the situation of the previous section, the process given by Algorithm 2 converges to the Brownian fan in $\mathcal{C}([0, T], \ell^p(\mathcal{M}))$ for any $T > 0$ and $p \leq 1$. The overall strategy of the proof is classical: we first prove a tightness result in Section 6.1 and then show that finite-dimensional marginals converge to those of the Brownian fan in Section 6.2.

Difficulties arise on two fronts. First, to prove the tightness result, it is convenient to have uniform moment bounds on the number of particles at fixed time for the approximating system. These turn out to be much more difficult to obtain for the approximating system than for the Brownian fan, which is mainly due to a lack of uniform exponential bounds. A second difficulty arises in the proof that finite-dimensional distributions converge to those of the Brownian fan. While it is intuitively clear that those excursions that survive for times of order $\mathcal{O}(1)$ do converge to suitably normalised Brownian excursions, this result is rather technical and, surprisingly, does not seem to appear in the literature. Furthermore, no convergence result holds for the typical excursions which die very early. We therefore also need to argue that, both at the level of Algorithm 2 and at the level of the Brownian fan, these small excursions do not matter in the limit.

6.1 Tightness

As in the previous section, we restrict ourselves to the particular case when the underlying Markov process is given by a rescaled random walk, namely

$$y_{(k+1)\varepsilon} = y_{k\varepsilon} + \sqrt{\varepsilon} \xi_{k+1}, \quad (6.1)$$

where the ξ_k are i.i.d. random variables with distribution ν having some exponential moments, and where the potential V is given by a linear function, $V(x) = -ax$. Our aim is to show that as $\varepsilon \rightarrow 0$, the sequence of birth and death processes obtained by running Algorithm 2 is tight in a state space \mathcal{X} , which we will now describe.

6.1.1 Formulation of the tightness result

With the construction of the previous section in mind, we choose as our state space $\mathcal{X} = \ell^p(\mathcal{M})$, where $\mathcal{M} = \{(x, v, n) \in \mathbf{R}^2 \times \mathbf{N} : v > -ax\}$, and $p \leq 1$ is arbitrary. Here, the coordinate n is used to keep track of the generation of a particle: direct offspring of a particle from the n th generation belong to the $(n + 1)$ st generation. We extend the Euclidean distance to $\mathbf{R}^2 \times \mathbf{N}$ by additionally postulating that the distance between particles belonging to different generations is given by the sums of the distances of the two particles to $\partial\mathcal{M}$. The boundary $\partial\mathcal{M}$ is given as before by $\partial\mathcal{M} = \{(x, v, n) : v = -ax\}$. We will use capital letters for elements of \mathcal{M} to differentiate them from elements of \mathbf{R} .

In order to formulate our result, we will make use of the following notation. For $t = k\varepsilon$ with k an integer, we denote by μ_t^ε the empirical measure of the particles alive at time t , and by N_t the number of such particles. Sometimes, it will be convenient to consider the particles instead as a collection of elements in $X_t^{(j)} \in \mathcal{M}$, so that we write

$$\mu_t^\varepsilon = \sum_{j=1}^{N_t} \delta_{X_t^{(j)}}.$$

We do not specify how exactly we order the particles, as this is completely irrelevant for our purpose. For $t \in (k\varepsilon, (k+1)\varepsilon)$, we define μ_t^ε by using the “linear interpolation” procedure (5.36), setting

$$\mu_t^\varepsilon = L_s(\mu_{k\varepsilon}^\varepsilon, \mu_{(k+1)\varepsilon}^\varepsilon), \quad s = \varepsilon^{-1}t - k.$$

The interpolation procedure L_s is a very minor modification from the one described above, in the sense that we only connect particles belonging to the same generation. In this way, the process $t \mapsto \mu_t^\varepsilon$ has continuous trajectories for every ε . The main result of this section is as follows:

Theorem 6.1 *Let $p \leq 1$ and denote by \mathcal{L}^ε the law of the process $t \mapsto \mu_t^\varepsilon$ described above, viewed as a family of probability measures on $\mathcal{C}([0, 1], \mathcal{X})$. Assume furthermore that there exists $c > 0$ such that $\int e^{c|y|} \nu(dy) < \infty$.*

Then, for any single particle initial condition $\mu_0^\varepsilon = \delta_{X_0}$ with $X_0 \in \mathcal{M}$, there exists $\varepsilon_0 > 0$ such that the family $\{\mathcal{L}^\varepsilon\}_{\varepsilon \leq \varepsilon_0}$ is tight.

Proof. Combining [Daw93, Theorem 3.6.4] and [Bil99, Theorem 8.3], we see that, in order to obtain tightness, it is sufficient to show that:

- For every $\delta > 0$, there exists a compact set $K_\delta \subset \mathcal{X}$ such that $\mathbf{P}(\mu_t^\varepsilon \in K_\delta) > 1 - \delta$, uniformly over $t \in [0, 1]$ and $\varepsilon < \varepsilon_0$.
- There exists $\alpha > 0$ and $C > 0$ such that $\mathbf{E}\|\mu_t^\varepsilon - \mu_s^\varepsilon\|_p^q \leq C|t - s|^{\alpha q}$, uniformly over all $s, t \in [0, 1]$ and $\varepsilon < \varepsilon_0$.

The first claim then follows from Proposition 6.7 below, while the second claim is the content of Proposition 6.8. \square

The proof of this result is the content of the remainder of this subsection and goes roughly as follows. In Section 6.1.2, we obtain a moment bound on the number of particles alive at any fixed time $t \in [0, 1]$, which is uniform in $\varepsilon > 0$. This then allows to obtain the compactness at fixed time in Section 6.1.3. The verification of Kolmogorov's continuity criterion is the content of Section 6.1.4.

6.1.2 Moment bounds on the number of particles

In the sequel, we will denote by $\mathbf{K}_p(X)$ the p th cumulant of a random variable X and by $\mathbf{K}_p(X | \mathcal{F})$ the same cumulant, conditioned on the σ -field \mathcal{F} . We will use the important property that $\mathbf{K}_p(X + Y | \mathcal{F}) = \mathbf{K}_p(X | \mathcal{F}) + \mathbf{K}_p(Y | \mathcal{F})$, provided that X and Y are independent, conditionally on \mathcal{F} . We will also use the fact that if X is a positive random variable, then $\mathbf{K}_p(X) \leq \mathbf{E}X^p$ and there exists a constant C such that the bound

$$\mathbf{E}X^p \leq C \sum_{q \leq p} (\mathbf{K}_q(X))^{p/q}, \quad (6.2)$$

holds.

Our aim now is to obtain a bound on the cumulants of the number of particles alive at time t which is independent of ε . We start with an initial configuration containing only one particle, which belongs to generation 0, and we set $N_t^0 \in \{0, 1\}$, depending on whether or not this particle is still alive at some subsequent time t . We also define $N_t^n = \mu_t^\varepsilon(\mathcal{M}^n)$, where $\mathcal{M}^n = \{(x, v, \ell) \in \mathcal{M} : \ell = n\}$, which is the number of particles in the n th generation that are alive at time t . We also denote by $N_{s,t}^n$ the number of such particles that were created at time $s \leq t$.

For any $X_0 = (x, v)$ with $v > -ax$, we write \mathbf{E}_{X_0} for expectations of observables for the process generated by starting Algorithm 2 with underlying dynamic (6.1), started with one single initial particle in generation 0 at location x with tag v . Finally, for $x \in \mathbf{R}$, we write \mathbf{E}_x for the same expectation, but where the initial particle has tag $v = +\infty$, meaning that it is “immortal”. We then have the following result:

Proposition 6.2 *Consider the situation of Theorem 6.1. For every $p \geq 0$, there exist $C_p > 0$ and $\varepsilon_0 > 0$ such that, under the rules of Algorithm 2, the number $N_t = \mu_t^\varepsilon(\mathcal{M})$ of particles alive at time t satisfies $\mathbf{E}_x |N_t|^p \leq 2 \exp(C_p t)$, uniformly over all $\varepsilon \leq \varepsilon_0$ and $x \in \mathbf{R}$. Furthermore, there exists $\varrho > 0$ such that*

$$\mathbf{E}_x N_t^n \leq (n+1)^{t/\varrho} 2^{-n}, \quad (6.3)$$

uniformly over all $\varepsilon \leq \varepsilon_0$, $n \geq 0$, and $t > 0$.

Finally, denoting by R_t^γ the number of offspring alive at time t that have never been at distance more than γ from $\partial\mathcal{M}$, we have the bound

$$\mathbf{E}_x R_t^\gamma \leq C\gamma,$$

uniformly over $t \in [0, 1]$, $\gamma \in (0, 1]$, and $\varepsilon \leq \varepsilon_0 \wedge \gamma^2$.

The proof will make use of the following elementary fact where, for $\lambda = n + p$ with $n \in \mathbf{N}$ and $p \in (0, 1]$, we denote by $\mathcal{I}(\lambda)$ the law of a random variable Y such that $Y = n$ with probability $1 - p$ and $Y = n + 1$ with probability p . (This is so that $\mathbf{E}Y = \lambda$.)

Lemma 6.3 *Let Y be a random variable with law $\mathcal{I}(\lambda)$. Then, for any $q \geq 1$ and any $\lambda > 0$, one has the bound $\mathbf{E}Y^q \leq \lambda + (2\lambda)^q$.*

Proof. By inspection, one has

$$\mathbf{E}Y^q = p(n+1)^q + (1-p)n^q.$$

In the case $n = 0$, one then has $\mathbf{E}Y^q = \lambda$, so the statement is true. For $n \geq 1$, one uses the fact that both $n+1$ and n are bounded by 2λ , and the claim follows at once. \square

Proof of Proposition 6.2. We restrict ourselves to times that are integer multiples of ε . Furthermore, from now on, we fix an initial condition x , so that we just write \mathbf{E} instead of \mathbf{E}_x . We also denote by \mathbf{F}^n the σ -algebra containing all information pertaining to particles in generations up to (and including) n . With this notation at hand, we obtain for N_t^n the bound

$$\begin{aligned} \mathbf{E}(N_t^n)^p &\lesssim \mathbf{E} \sum_{q=1}^p (\mathbf{K}_q(N_t^n | \mathbf{F}^{n-1}))^{p/q} \\ &= \mathbf{E} \sum_{q=1}^p \left(\sum_{\varepsilon\ell \leq t} \mathbf{K}_q(N_{\varepsilon\ell,t}^n | \mathbf{F}^{n-1}) \right)^{p/q} \\ &\lesssim \sum_{q=1}^p \mathbf{E} \left(\sum_{\varepsilon\ell \leq t} \mathbf{E}(|N_{\varepsilon\ell,t}^n|^q | \mathbf{F}^{n-1}) \right)^{p/q}, \end{aligned} \quad (6.4)$$

where we used (6.2) in the first step and the independence of the offspring in the second step. In the above expression, ℓ takes only integer values. Note now that

$$\sum_{\varepsilon\ell \leq t} \frac{\varepsilon}{\sqrt{t(t+\varepsilon-\varepsilon\ell)}} \leq C, \quad (6.5)$$

uniformly over all $t \geq \varepsilon$. As a consequence, for any positive sequence a_n and any power $r \geq 1$, one has the bound

$$\begin{aligned} \left(\sum_{\varepsilon\ell \leq t} a_\ell \right)^r &= \left(\sum_{\varepsilon\ell \leq t} \frac{\varepsilon}{\sqrt{t(t+\varepsilon-\varepsilon\ell)}} \tilde{a}_\ell \right)^r \lesssim \sum_{\varepsilon\ell \leq t} \frac{\varepsilon}{\sqrt{t(t+\varepsilon-\varepsilon\ell)}} \tilde{a}_\ell^r \\ &= \sum_{\varepsilon\ell \leq t} \frac{t^{\frac{r-1}{2}} (t+\varepsilon-\varepsilon\ell)^{\frac{r-1}{2}}}{\varepsilon^{r-1}} a_\ell^r, \end{aligned}$$

where we have set $\tilde{a}_\ell = \varepsilon^{-1} \sqrt{t(t+\varepsilon-\varepsilon\ell)} a_\ell$ in the intermediate steps. Applying this inequality to (6.4), we obtain

$$\mathbf{E}(N_t^n)^p \lesssim \sum_{q=1}^p \sum_{\varepsilon\ell \leq t} \frac{t^{\frac{p-q}{2q}} (t+\varepsilon-\varepsilon\ell)^{\frac{p-q}{2q}}}{\varepsilon^{\frac{p-q}{q}}} \mathbf{E} \left(\mathbf{E}(|N_{\varepsilon\ell,t}^n|^q | \mathbf{F}^{n-1}) \right)^{p/q}. \quad (6.6)$$

Denote now by $M_s^{n,j}$ the number of particles in the n th generation created at time s by the j th particle from the $n-1$ st generation. We write \mathbf{G}_s for the σ -algebra generated by this additional data. Each of these particles yields a contribution to $N_{s,t}^n$ of either 1 or 0, depending whether it survives or not. Furthermore, these contributions, which we will denote by $S_{s,t}^{n,j,i}$, are all independent and, for the same value of j , they are also identically distributed. By definition, we thus have the identity

$$N_{\varepsilon\ell,t}^n = \sum_{j=1}^{N_{\varepsilon\ell,t}^{n-1}} \sum_{i=1}^{M_{\varepsilon\ell}^{n,j}} S_{\varepsilon\ell,t}^{n,j,i}.$$

We now twice make use of the inequality

$$\left(\sum_{j=1}^m a_j\right)^q \leq m^{q-1} \sum_{j=1}^m a_j^q, \quad (6.7)$$

which is valid for any $q \geq 1$, $m \geq 0$ and sequence of positive numbers a_j . This yields the bound

$$(N_{\varepsilon\ell,t}^n)^q \leq (N_{(\ell-1)\varepsilon}^{n-1})^{q-1} \sum_{j=1}^{N_{(\ell-1)\varepsilon}^{n-1}} (M_{\varepsilon\ell}^{n,j})^{q-1} \sum_{i=1}^{M_{\varepsilon\ell}^{n,j}} S_{\varepsilon\ell,t}^{n,j,i}. \quad (6.8)$$

Note that since $S_{\varepsilon\ell,t}^{n,j,i}$ can only take the values 0 or 1, raising it to the power q makes no difference. By the “gambler’s ruin theorem” [LL10, Thm 5.1.7], we have the bound

$$\mathbf{E}(S_{\varepsilon\ell,t}^{n,j,i} | \mathbf{F}^{n-1} \vee \mathbf{G}_{\varepsilon\ell}) \leq C\sqrt{\varepsilon} \frac{\xi^j + 1}{\sqrt{t + \varepsilon - \ell\varepsilon}}, \quad (6.9)$$

where $\sqrt{\varepsilon}\xi^j$ denotes the step performed by the j th particle of the $(n-1)$ st generation between times $\varepsilon(\ell-1)$ and $\varepsilon\ell$. Regarding the number of offspring $M_{\varepsilon\ell}^{n,j}$, it follows from the definition of the algorithm that its distribution is given by $\mathcal{I}(\exp(a\sqrt{\varepsilon}\xi_+^j) - 1)$, where ξ_+^j denotes the positive part of ξ^j . Combining this with (6.8), Lemma 6.3, and (6.9), we thus obtain the bound

$$\mathbf{E}(|N_{\varepsilon\ell,t}^n|^q | \mathbf{F}^{n-1}) \lesssim |N_{(\ell-1)\varepsilon}^{n-1}|^{q-1} \sum_{j=1}^{N_{(\ell-1)\varepsilon}^{n-1}} \sqrt{\varepsilon} \frac{(e^{a\sqrt{\varepsilon}\xi_+^j} - 1) + (e^{a\sqrt{\varepsilon}\xi_+^j} - 1)^q}{\sqrt{t + \varepsilon - \ell\varepsilon}} (\xi^j + 1).$$

In order to simplify this expression, we use the fact that, for $x \geq 0$, there exists a constant C depending on q such that

$$e^x - 1 \leq xe^x, \quad (e^x - 1)^q \leq Cxe^{qx},$$

for every $q \geq 1$. This yields

$$\mathbf{E}(|N_{\varepsilon\ell,t}^n|^q | \mathbf{F}^{n-1}) \lesssim |N_{(\ell-1)\varepsilon}^{n-1}|^{q-1} \sum_{j=1}^{N_{(\ell-1)\varepsilon}^{n-1}} \varepsilon \frac{\mathbf{1}_{\xi^j \geq 0} e^{aq\sqrt{\varepsilon}\xi^j}}{\sqrt{t + \varepsilon - \ell\varepsilon}} (\xi^j + 1)^2.$$

Using (6.7) once again, we get the bound

$$(\mathbf{E}(|N_{\varepsilon\ell,t}^n|^q | \mathbf{F}^{n-1}))^{\frac{p}{q}} \lesssim |N_{(\ell-1)\varepsilon}^{n-1}|^{p-1} \sum_{j=1}^{N_{(\ell-1)\varepsilon}^{n-1}} \varepsilon^{\frac{p}{q}} \frac{\mathbf{1}_{\xi^j \geq 0} e^{ap\sqrt{\varepsilon}\xi^j}}{(t + \varepsilon - \ell\varepsilon)^{\frac{p}{2q}}} (\xi^j + 1)^{\frac{2p}{q}}.$$

At this stage, we note that, conditional on the state of the system at time $\varepsilon(\ell-1)$ the steps ξ^j are all independent and identically distributed with law ν . Setting

$$P_\varepsilon = \int_0^\infty e^{ap\sqrt{\varepsilon}z} (z+1)^{\frac{2p}{q}} \nu(dz),$$

it follows that

$$\mathbf{E}(\mathbf{E}(|N_{\varepsilon\ell,t}^n|^q | \mathbf{F}^{n-1}))^{\frac{p}{q}} \lesssim \frac{\varepsilon^{\frac{p}{q}} P_\varepsilon}{(t + \varepsilon - \ell\varepsilon)^{\frac{p}{2q}}} \mathbf{E}|N_{(\ell-1)\varepsilon}^{n-1}|^p.$$

Inserting this into (6.6) yields for $t \leq 1$ the bound

$$\mathbf{E}(N_t^n)^p \lesssim P_\varepsilon \sum_{\varepsilon \ell \leq t} \frac{\varepsilon}{\sqrt{t + \varepsilon - \varepsilon \ell}} \mathbf{E}|N_{(\ell-1)\varepsilon}^{n-1}|^p.$$

On the other hand, since we assumed that the initial particle is immortal, N_s^{n-1} is stochastically increasing as a function of s , so that we obtain from (6.5) the bound

$$\mathbf{E}(N_t^n)^p \lesssim P_\varepsilon \sqrt{t} \mathbf{E}(N_t^{n-1})^p.$$

Since, by the exponential integrability assumption on ν , there exists $C > 0$ such that, for ε small enough, $P_\varepsilon \leq C$ and since $N_t^0 = 1$, we conclude that there exists some λ_p such that, for all $n \geq 1$, one has the bound

$$\mathbf{E}(N_t^n)^p \lesssim (\lambda_p t)^{n/2}. \quad (6.10)$$

It follows that there exists $t_* > 0$ (depending on p) such that $\mathbf{E}(N_t)^p \leq 2$ (say), for every $t \leq t_*$. To show that $\mathbf{E}(N_t)^p < \infty$ for every t , we use again (6.7), combined with the Markov property of the process to conclude that $\mathbf{E}(N_{t+t_*})^p \leq 2\mathbf{E}(N_t)^p$, from which the first claim follows.

To get (6.3), observe that if we choose ϱ small enough so that $\lambda_1 \varrho \leq 1/4$, the bound with $t \leq \varrho$ follows from (6.10). Denote now by \mathcal{F}_t the filtration generated by all events up to time t . For t and ϱ that are multiples of ε it follows from the Markov property that

$$\mathbf{E}(N_t^n | \mathcal{F}_{t-\varrho}) \leq \sum_{\ell=0}^n N_{t-\varrho}^\ell \mathbf{E} N_\varrho^{n-\ell}$$

where the expectation on the right is taken with respect to an initial condition with one single immortal particle. The claimed bound for arbitrary $t > 0$ therefore follows by induction.

It remains to obtain the bound on R_t^γ . Denote by $R_{s,t}^\gamma$ the number of offspring contributing to R_t^γ that are created at time $s < t$, so that $R_t^\gamma = \sum_{\varepsilon k \leq t} R_{k\varepsilon,t}^\gamma$. Denote now by $Q_{z,t}^\gamma$ the probability that, after time t , the random walk (6.1) with initial condition $\sqrt{\varepsilon}z$ has never exited the interval $[0, \delta]$. It then follows that

$$\mathbf{E}(R_{k\varepsilon,t}^\gamma | \mathcal{F}_{k\varepsilon}) = a\sqrt{\varepsilon} N_{k\varepsilon} \int_0^\infty \int_0^z e^{a\sqrt{\varepsilon}y} Q_{y,t-k\varepsilon}^\gamma dy \nu(dz). \quad (6.11)$$

This is because, by Step 3 of Algorithm 2, the expected number of offspring created by a particle performing an upward step of size $\sqrt{\varepsilon}z$ is given by $e^{a\sqrt{\varepsilon}z} - 1$ while, if we denote by $\sqrt{\varepsilon}y$ the distance between the starting point of the offspring and the “wall” below which it is killed, the law of y is given by $c^{-1} \int_0^z e^{a\sqrt{\varepsilon}y} dy$ where c is a normalizing constant. Since $c = (e^{a\sqrt{\varepsilon}z} - 1)/a\sqrt{\varepsilon}$, (6.11) follows.

Using again the gambler’s ruin theorem, combined with the Markov property of the random walk, we obtain for $Q_{z,t}^\gamma$ the bound

$$Q_{z,t}^\gamma \leq C \frac{(z+1)\sqrt{\varepsilon}}{\sqrt{t+\varepsilon}} \bar{Q}_{t/2}^\gamma,$$

where \bar{Q}_t^γ is the probability that the random walk (6.1), starting at the origin, stays within $[-\gamma, \gamma]$ up to time t . Using the scaling properties of Brownian motion, combined with [Fra73], we conclude that one has the bound

$$\bar{Q}_t^\gamma \leq 1 \wedge \frac{C\delta^{2q}}{t^q},$$

for every $q > 0$, so that

$$Q_{z,t}^\gamma \lesssim \frac{(z+1)\sqrt{\varepsilon}}{\sqrt{t+\varepsilon}} \left(1 \wedge \frac{\delta^{2q}}{t^q}\right).$$

Inserting this bound into (6.11), using the previously obtained bounds on $N_{k\varepsilon}$, and summing over all k , the claim follows at once. \square

Remark 6.4 It follows from Theorem 4.1 that, in the particular case when the initial tag v is distributed according to the logarithm of a uniformly distributed random variable, one has the identity $\mathbf{E}N_t = \mathbf{E}e^{ay_t}$, where y_t is the rescaled random walk with steps ν . The (at least one-sided) exponential integrability assumption on ν is therefore a necessary assumption in order to obtain any kind of moment bounds on N_t .

Remark 6.5 Even if we assume that ν has Gaussian tails and despite the result previously obtained in Theorem 5.13, it is *not* true in general that N_t has uniform exponential moments as $\varepsilon \rightarrow 0$. This is because even for the first step, the probability that the original particle performs a step of order ε^{-p} is of order $\exp(-\varepsilon^{-2p-1})$. If this were to happen, the number of offspring created in this way would be of order $\exp(\varepsilon^{-p})$, which immediately shows that exponential moments blow up as $\varepsilon \rightarrow 0$.

6.1.3 Compactness at any fixed time

We now show that, for any fixed time $t \in [0, 1]$, we can find a compact set K_δ such that the process μ_t^ε belongs to K_δ with high probability, uniformly over ε . Our first ingredient for this is the following moment bound:

Proposition 6.6 *Consider the setting of Theorem 6.1. Then, there exist constants C and ε_0 such that, for every $t \in [0, 1]$ and every $\varepsilon \leq \varepsilon_0$, the bound*

$$\mathbf{E}_{X_0} \left(\int_{\mathcal{M}} |x - x_0|^{2p} \mu_t^\varepsilon(dx, dv, dk) \right) \leq Ct^p, \quad (6.12)$$

holds uniformly over all initial conditions $X_0 = (x_0, v_0, 0) \in \mathcal{M}$.

Proof. We restrict ourselves to the case when t is an integer multiple of ε , since the bound on the remaining times easily follows from our interpolation procedure. Furthermore, we can restrict ourselves to the case when the initial particle is immortal, which formally corresponds to setting $v = +\infty$. By translation invariance, we also restrict ourselves to the case where the initial particle is located at the origin, and we denote the corresponding expectation by \mathbf{E}_0 .

It follows from Theorem 4.1 that if we choose $v = \frac{1}{a} \log u$, where u is drawn uniformly from $[0, 1]$ and denote by $\tilde{\mu}_t^\varepsilon$ the corresponding process, then one has the identity

$$\mathbf{E}_0 \left(\int_{\mathcal{M}} x^{2p} \tilde{\mu}_t^\varepsilon(dx, dv, dk) \right) = \mathbf{E}_0 \left(e^{ay_t} (y_t)^{2p} \right),$$

where y_t is the simple random walk (6.1) started at the origin. It follows immediately from the exponential integrability of ν that this quantity is bounded by Ct^p for $t \leq 1$ and for ε small enough.

On the other hand, one can realise the process μ_t^ε , which starts with an immortal initial particle, in the following way:

1. Consider the process $\tilde{\mu}_t^\varepsilon$, where the v -component is as above.
2. When the initial (generation 0) particle is killed, replace it instantly by an immortal particle starting at the current location.

Let x_t^0 denote the trajectory of the initial particle and let s be the time at which the initial particle is killed and replaced. Let $P_t = \mathbf{P}(s \leq t)$. By the construction just outlined, we have the recursion relation

$$F_{t,0} = E_t + P_t \mathbf{E}(F_{t-s, x_s^0} \mid s \leq t) ,$$

where we set

$$\begin{aligned} F_{t, x_s^0} &= \mathbf{E}_0 \left(\int_{\mathcal{M}} (x + x_s^0)^{2p} \mu_t^\varepsilon(dx, dv, dk) \right) , \\ E_t &= \mathbf{E}_0 \left(\int_{\mathcal{M}} x^{2p} \tilde{\mu}_t^\varepsilon(dx, dv, dk) \right) . \end{aligned}$$

(Remember that the difference between E_t and $F_{t,0}$ is that in order to compute F , we start with an immortal particle.) Setting $F_t = F_{t,0}$, using the fact that, for every $\delta > 0$ one can find C_δ such that $(x + x_s^0)^{2p} \leq C_\delta (x_s^0)^{2p} + (1 + \delta)x^{2p}$, and recalling that $E_t \leq Ct^p$, we deduce that

$$F_t \leq Ct^p + C_\delta P_t \mathbf{E}((x_s^0)^{2p} \mathbf{E}_{x_s^0}(N_{t-s}) \mid s \leq t) + (1 + \delta)P_t \mathbf{E}(F_{t-s} \mid s \leq t) ,$$

where N_t is the number of particles alive at time t for the system started with an immortal particle. Note now that, for $t \leq 1$, we know from Proposition 6.2 that the expected number of particles alive at any given time is bounded by some constant uniform in ε . Since this bound is also uniform in the initial condition, we have

$$P_t \mathbf{E}((x_s^0)^{2p} \mathbf{E}_{x_s^0}(N_{t-s} \mid s \leq t)) \lesssim P_t \mathbf{E}((x_s^0)^{2p} \mid s \leq t) = \mathbf{E}((x_s^0)^{2p} \mathbf{1}_{s \leq t}) .$$

Defining $\tilde{s} = s \wedge t$, we then obtain the trivial bound

$$\mathbf{E}((x_s^0)^{2p} \mathbf{1}_{s \leq t}) \leq \mathbf{E}_0 |y_{\tilde{s}}|^{2p} \leq \mathbf{E}_0 |y_t|^{2p} \lesssim t^p ,$$

where we made use of the fact that $|y_t|^{2p}$ is a submartingale to obtain the second inequality.

Setting now $\bar{F}_t = \sup_{s \leq t} F_s$, we can combine all of these bounds to get the inequality

$$\bar{F}_t \leq C_\delta t^p + (1 + \delta)P_t \bar{F}_t .$$

Since one can check (for example by again using the fact that the random walk approximates a Brownian motion for small ε) that $\sup_{\varepsilon \leq 1} \sup_{t \leq 1} P_t = \sup_{\varepsilon \leq 1} P_1 < 1$, the claim follows at once by choosing δ sufficiently small. \square

This result can now be used to deduce the announced uniform tightness result over fixed times:

Proposition 6.7 *Consider the setting of Theorem 6.1. Then, for every $\delta > 0$ there exists a compact set $K_\delta \subset \mathcal{X}$ such that $\mathbf{P}_{X_0}(\mu_t^\varepsilon \in K_\delta) > 1 - \delta$, uniformly over $t \in [0, 1]$, $\varepsilon \leq \varepsilon_0$ and $X_0 \in \mathcal{M}^0$.*

Proof. For any $n \in \mathbf{Z}_+$ and $R \in \mathbf{R}_+$, denote

$$\mathcal{M}^{[n]} = \{(x, v, \ell) \in \mathcal{M} : \ell \leq n\}, \quad B_R = \{(x, v, \ell) \in \mathcal{M} : |x| \vee |v| \leq R\}.$$

For any such n and R and for $m \in \mathbf{Z}_+$, we then denote by $K_{n,m,R} \subset \mathcal{X}$ the set of all integer-valued measures η on \mathcal{M} such that $\eta(\mathcal{M} \setminus \mathcal{M}^{[n]}) = 0$, $\eta(\mathcal{M}) \leq m$, and $\eta(\mathcal{M} \setminus B_R) = 0$. Since these sets are obviously compact, it suffices to find, for every $\delta > 0$, sufficiently large values n , R and m so that $\mathbf{P}_{X_0}(\mu_t^\varepsilon \in K_{n,m,R}) > 1 - \delta$ uniformly over all $\varepsilon \leq \varepsilon_0$ and $t \in [0, 1]$.

Note that $K_{n,m,R} \subset K_n^1 \cap K_m^2 \cap K_R^x \cap K_R^v$ where, for K_n^1 and K_m^2 we only enforce the conditions involving n and m respectively. The set K_R^x consists of configurations such that the x -coordinate of every particle is less than R in absolute value, while K_R^v enforces that the v -coordinate be less than R .

It follows immediately from (6.3) that there exists $\gamma > 0$ and a constant C (depending in principle on the time we consider, but it can be chosen uniformly over $t \in [0, 1]$), such that

$$\mathbf{P}_{X_0}(\mu_t^\varepsilon \notin K_n^1) \leq Ce^{-\gamma n},$$

for every $n \geq 0$. Similarly, it follows from the moment bounds on N_t obtained in Proposition 6.2 that, for every $p > 0$, there exists C such that

$$\mathbf{P}_{X_0}(\mu_t^\varepsilon \notin K_m^2) \leq \frac{C}{m^p}.$$

In order to get a bound on the x -coordinate of the particles, we combine Proposition 6.6 with Chebyshev's inequality, so that

$$\mathbf{P}_{X_0}(\mu_t^\varepsilon \notin K_R^x) \leq \frac{C}{R^2}.$$

It remains to obtain a bound on the probability of not being in K_R^v . For this, we use the fact that on the one hand, the label of a particle always satisfies $v > -ax$. On the other hand, any descendant of the initial particle will always satisfy $v \leq -a \sup_{t \leq 1} x_t^0$, where here we denote by x_t^0 the position of the original particle at time t . Since this particle was assumed to start at the origin, we obtain $v^2 \leq a^2 x^2 + a^2 \sup_{t \leq 1} (x_t^0)^2$, so that we obtain the bound

$$\mathbf{P}_{X_0}(\mu_t^\varepsilon \notin K_R^v) \leq \frac{C}{R^2},$$

just as above. Combining all of these bounds, the claim follows by choosing n , R and m large enough. \square

6.1.4 Kolmogorov criterion

The aim of this section is to obtain the following bound on the time regularity of our process:

Proposition 6.8 *For every $p \leq 1$ and $q \geq 1$, there exists a constant C such that*

$$\mathbf{E}_{X_0} \|\mu_\delta^\varepsilon - \delta_{X_0}\|_p^q \leq C \delta^{pq/2}, \quad (6.13)$$

where $X_0 \in \mathcal{M}^0$ is an initial condition with only one particle in the system, and $\delta < 1$ is an integer multiple of ε .

Remark 6.9 As usual, the precise location of the particle is irrelevant by translation invariance, so the above bound is uniform over all choices of X_0 .

Before we proceed to the proof of Proposition 6.8, we observe that the bound (6.13) implies that Kolmogorov's criterion holds for the process over a fixed interval of time:

Corollary 6.10 *For every $p \leq 1$ and $q \geq 1$, one has the bound*

$$\mathbf{E}_{X_0} \|\mu_{t+\delta}^\varepsilon - \mu_t^\varepsilon\|_p^q \leq C \delta^{pq/2}, \quad (6.14)$$

uniformly over all $\delta, t, \varepsilon \in [0, 1]$.

Proof. Note first that we can restrict ourselves to the case when t and δ are integer multiples of ε . Indeed, it follows from the definition of $\|\cdot\|_p$ and from (5.37) that, if $k\varepsilon \leq s < t \leq (k+1)\varepsilon$, then

$$\|\mu_s^\varepsilon - \mu_t^\varepsilon\| \leq \varepsilon^{-p} |t - s|^p \|\mu_{k\varepsilon}^\varepsilon - \mu_{(k+1)\varepsilon}^\varepsilon\|.$$

We then obtain from Proposition 6.8 the bound

$$\begin{aligned} \mathbf{E}_{X_0} \|\mu_{t+\delta}^\varepsilon - \mu_t^\varepsilon\|_p^q &= \mathbf{E}_{X_0} \left(\mathbf{E}_{X_0} (\|\mu_{t+\delta}^\varepsilon - \mu_t^\varepsilon\|_p^q | \mathcal{F}_t) \right) \\ &\leq \mathbf{E}_{X_0} \left(|N_t|^{q-1} \sum_{j=1}^{N_t} \mathbf{E}_{X_t^{(j)}} \|\mu_\delta^\varepsilon - \delta_{X_t^{(j)}}\|_p^q \right) \\ &\leq C \delta^{pq/2} \mathbf{E}_{X_0} (|N_t|^q) \leq C \delta^{pq/2}, \end{aligned}$$

where we made use of the Markov property, (6.7), and (6.14). \square

Proof of Proposition 6.8. Denote by X_t the location at time t of a single particle starting at $X_0 = (x_0, v_0, 0)$ and evolving under the rescaled random walk stopped when it reaches the boundary of \mathcal{M} . Denote by x_t the position in \mathbf{R} corresponding to X_t . From the properties of the random walk and the definition of $\|\cdot\|_p$, for every $q > 1$ there exists a constant C such that the bound

$$\mathbf{E}_{X_0} \|\delta_{X_\delta} - \delta_{X_0}\|_p^q \leq C \delta^{qp/2},$$

holds, independently of the initial condition and independently of $\varepsilon \leq 1$.

Let us now bound the contribution from the descendants of the initial particle. Ordering the particles alive at time t in such a way that the original particle has label 1 (if it is not alive anymore, we consider it as being located on the boundary, where it was stopped), we have the bound

$$\|\delta_{X_\delta} - \mu_\delta^\varepsilon\|_p \leq \sum_{j=2}^{N_\delta} d_p(X_\delta^{(j)}, \partial\mathcal{M}),$$

so that

$$\begin{aligned} \mathbf{E}_{X_0} \|\delta_{X_\delta} - \mu_\delta^\varepsilon\|_p^q &\leq \mathbf{E}_{X_0} \left(\sum_{j=2}^{N_\delta} d_p(X_\delta^{(j)}, \partial\mathcal{M}) \right)^q \\ &\leq \left(\left(\mathbf{E}_{X_0} (N_\delta - 1)^{2q-1} \right) \left(\mathbf{E}_{X_0} \sum_{j=2}^{N_\delta} d_p^{2q}(X_\delta^{(j)}, \partial\mathcal{M}) \right) \right)^{1/2} \end{aligned}$$

$$\leq C \left(\mathbf{E}_{X_0} \sum_{j=2}^{N_\delta} d_p^{2q}(X_\delta^{(j)}, \partial\mathcal{M}) \right)^{1/2},$$

where the second inequality follows from Proposition 6.2. Recall now that if $X_\delta^{(j)} = (x, v, n)$, then one has $d_p(X_\delta^{(j)}, \partial\mathcal{M}) = |x - v/a|^p$, where v/a is guaranteed to take values between $\inf_{s \leq \delta} x_s$ and x . As a consequence, we have the bound

$$|x - v/a|^p \leq |x - \inf_{s \leq \delta} x_s|^p \leq |x - x_0|^p + \sup_{s \leq \delta} |x_s - x_0|^p,$$

so that

$$\begin{aligned} \mathbf{E}_{X_0} \sum_{j=2}^{N_\delta} d_p^{2q}(X_\delta^{(j)}, \partial\mathcal{M}) &\lesssim \mathbf{E}_{X_0} \left(N_\delta \sup_{s \leq \delta} |x_s - x_0|^{2pq} \right) \\ &\quad + \mathbf{E}_{X_0} \int |x - x_0|^{2pq} \mu_\delta^\varepsilon(dx, dv, dk) \lesssim \delta^{pq}, \end{aligned}$$

where the last bound is a consequence of Propositions 6.2 and 6.6, as well as standard bounds on the supremum of a random walk. The claim now follows at once. \square

6.2 Convergence of fixed-time distributions

In this section, we show that any limiting process obtained from the tightness result of the previous section necessarily coincides with the Brownian fan constructed in Section 5.3.2. With the notations of that section at hand, our convergence result can be formulated as follows.

Theorem 6.11 *Consider the setting of Theorem 6.1 with an initial condition $X_0 = (x, v, 0) \in \mathcal{M}$, and consider μ_t as above with the “initial condition” for $\mu^{[\infty]}$ given by a Brownian motion starting at x , killed when it reaches the level $-v/a$.*

Then, for every $p \in (0, 1]$, there exists a version of the process $\{\mu_t\}_{t \geq 0}$ which is a continuous Markov process with values in $\ell^p(\mathcal{M})$. Furthermore, denoting the law of its restriction to the time interval $[0, 1]$ by \mathcal{L} , the sequence of measures \mathcal{L}^ε converges weakly to \mathcal{L} in $\mathcal{C}([0, 1], \ell^p(\mathcal{M}))$.

Proof. It suffices to show that, for any fixed collection of times $\{t_1, \dots, t_k\}$, the law of $\{\mu_{t_i}^\varepsilon\}_{i \leq k}$ converges weakly to that of $\{\mu_{t_i}\}_{i \leq k}$. Indeed, Corollary 6.10 then implies that the process μ_t satisfies Kolmogorov’s continuity criterion and therefore has a continuous version. By Theorem 6.1, we deduce weak convergence in $\mathcal{C}([0, 1], \ell^p(\mathcal{M}))$ from the convergence of marginals. Using the Markov property, the superposition property of the process, and Proposition 6.2, we reduce ourselves to the case $k = 1$ with an initial condition consisting of one single particle.

Denote now by $\mu^{\varepsilon, [\infty]}$ the random integer-valued measure on \mathcal{E} obtained by running Algorithm 2 until time 1. Observe that $\mu^{\varepsilon, [\infty]}$ can be built in the following way. For an excursion $w \in \mathcal{E}$, we build a random measure $\mathbf{Q}^\varepsilon(w)$ by the following procedure. For every $k \in \mathbf{N}$ with $\varepsilon k > \mathfrak{s}(w)$ and $\varepsilon(k+1) < \mathfrak{e}(w) \wedge 1$ we set

$$\Delta w_k = w_{(k+1)\varepsilon} - w_{k\varepsilon}.$$

If $\Delta w_k > 0$, we then draw a random variable N_k^1 with law $\mathcal{I}(\exp(a\Delta w_k) - 1)$ (see (2.9) and Lemma 6.3). For $j = 1, \dots, N_k^1$, we build i.i.d. excursions $w^{k,j} \in \mathcal{E}$ by the

following procedure. First, draw a uniform random variable $u \sim \mathcal{U}(\exp(-a\Delta w_k), 1)$ and set $v = \log u - aw_{k\varepsilon}$. Then, denote by $\{y_{\ell\varepsilon}^{k,j}\}_{\ell=0}^L$ an instance of the random walk (6.1), started at $w_{(k+1)\varepsilon}$ and stopped just before it becomes smaller than $-v/a$ (so that $y_{L\varepsilon}^{k,j} > -v/a$). The excursion $w^{k,j}$ is then given by $\mathfrak{s}(w^{k,j}) = k\varepsilon$, $\mathfrak{e}(w^{k,j}) = (L+k+2)\varepsilon$, $w_{\varepsilon(\ell+k+1)}^{k,j} = y_{\ell\varepsilon}^{k,j}$ for $\ell \in \{0, \dots, L\}$, $w^{k,j}(\ell\varepsilon) = -v/a$ for the remaining integer values of ℓ , and linear interpolation in between integer values. We then set

$$Q^\varepsilon(w) = \sum_{k=1}^{\infty} \sum_{j=1}^{N_k^1} \delta_{w^{k,j}},$$

which is the point measure describing the children of the particle with trajectory w . Similarly to before, we build $\mu^{\varepsilon, [\infty]}$ recursively by the following procedure:

- Build an excursion $w^0 \in \mathcal{E}$ as above, starting at x and stopped at $-v/a$, where $(x, v, 0) \in \mathcal{M}$ is the initial condition appearing in the statement. Set $\mu^{\varepsilon, 0} = \delta_{w^0}$.
- Given $\mu^{\varepsilon, \ell}$, define $\mu^{\varepsilon, \ell+1}$ by

$$\mu^{\varepsilon, \ell+1} = \int_{\mathcal{E}} Q^\varepsilon(w) \mu^{\varepsilon, \ell}(dw),$$

where the $\{Q^\varepsilon(w)\}_{w \in \mathcal{E}}$ are all independent (and independent of the $\mu^{\varepsilon, \ell'}$ with $\ell' \leq \ell$). Note that the integral is actually a finite sum, so the construction makes sense.

- Set $\mu^{\varepsilon, [\ell]} = \sum_{\ell'=0}^{\ell} \mu^{\varepsilon, \ell'}$ for positive ℓ (including the case $\ell = \infty$).

If we set

$$\mu_t^\varepsilon = E_t^* \mu^{\varepsilon, [\infty]},$$

where E_t is defined as in Section 5.3.2, then the process μ_t^ε is indeed equal in law to the process considered in Section 6.1¹.

Denote now by \mathcal{E}_γ the set of excursions of height at least γ , namely

$$\mathcal{E}_\gamma = \{w \in \mathcal{E} : \exists t \in \mathbf{I}(w) \text{ with } w_t \geq w(\mathfrak{s}(w)) + \gamma\}.$$

We also write $\mu_\gamma^{\varepsilon, [n]}$ for the measure obtained exactly like $\mu^{\varepsilon, [n]}$, but where we replace $Q^\varepsilon(w)$ by its restriction $Q_\gamma^\varepsilon(w)$ to the set \mathcal{E}_γ at every step, so that $\mu_\gamma^{\varepsilon, [n]} \leq \mu_\varepsilon^{[\infty]}$ almost surely. In words, Q_γ^ε is obtained from Q^ε by discarding all excursions of height less than γ , as well as the descendants of any such excursions.

Combining Propositions 6.2 and 6.6, we see that, for ε_0 small enough, there exist constants C and $\alpha > 0$ such that one has the bound

$$\sup_{\varepsilon \leq \varepsilon_0} \mathbf{P}(E_t^* \mu^{\varepsilon, [\infty]} \neq E_t^* \mu_\gamma^{\varepsilon, [n]}) \leq C(e^{-\alpha n} + \gamma^p), \quad (6.15)$$

uniformly over ε and $t \in [0, 1]$.

Following an argument along the lines of the proof of Theorem 5.13, a similar bound can be shown to hold for $\mathbf{P}(E_t^* \mu^{[\infty]} \neq E_t^* \mu_\gamma^{[n]})$, where $\mu_\gamma^{[n]}$ is the recursive Poisson process of depth n constructed like $\mu^{[\infty]}$, but with $Q(w, \cdot)$ replaced by its restriction

¹Strictly speaking, the two processes agree only at times that are multiples of ε since the two interpolation procedures we used may differ when the trajectories of two particles from the same generation cross each other. This is irrelevant.

$\mathcal{Q}_\gamma(w, \cdot)$ to \mathcal{E}_γ . As a consequence, we conclude that it is sufficient to show that, for every fixed $\gamma > 0$ and $n \geq 1$,

$$\lim_{\varepsilon \rightarrow 0} \mathcal{D}(\mu_\gamma^{\varepsilon, [n]}) = \mathcal{D}(\mu_\gamma^{[n]}), \quad (6.16)$$

where $\mathcal{D}(\cdot)$ denotes the law of a random variable and convergence is to be understood in the sense of weak convergence on the space $M_+(\mathcal{E})$ endowed with the Wasserstein-1 distance (see [Vil03] and Section 6.2.2 below).

Our aim then is to make use of Theorem 6.17 below which gives a general convergence result for recursive Poisson point processes. The drawback is that $\mu_\gamma^{\varepsilon, [n]}$ is itself not a recursive Poisson point process, due to the fact that $Q_\gamma^\varepsilon(w)$ is not a realisation of a Poisson point process. However, it would be one if, in the construction of $Q_\gamma^\varepsilon(w)$, we had drawn N_k^1 according to a Poisson distribution with mean $\exp(a\Delta w_k) - 1$. Denote by $\tilde{Q}_\gamma^\varepsilon(w)$ the Poisson point process obtained in this way, by $\tilde{Q}_\gamma^\varepsilon$ its restriction to \mathcal{E}_γ , and by $\tilde{\mu}_\gamma^{\varepsilon, [n]}$ the recursive Poisson point process of depth n obtained by iterating $\tilde{Q}_\gamma^\varepsilon$. We then claim that it is possible to find a coupling between $\tilde{\mu}_\gamma^{\varepsilon, [n]}$ and $\mu_\gamma^{\varepsilon, [n]}$ such that

$$\mathbf{P}(\tilde{\mu}_\gamma^{\varepsilon, [n]} \neq \mu_\gamma^{\varepsilon, [n]}) \leq C_{n, \gamma} \sqrt{\varepsilon}, \quad (6.17)$$

where the constant $C_{n, \gamma}$ depends on n and γ , but not on ε . Indeed, if we denote by U_λ a random variable with law $\mathcal{I}(\lambda)$ and by \bar{U}_λ a Poisson random variable with mean λ , then it is straightforward to check that there exists a constant C and a coupling between U_λ and \bar{U}_λ such that

$$\mathbf{P}(|U_\lambda - \bar{U}_\lambda| = 1) \leq C(1 \wedge \lambda^2), \quad \mathbf{P}(|U_\lambda - \bar{U}_\lambda| > 1) \leq C(1 \wedge \lambda^3).$$

Furthermore, by Proposition 5.1, the probability that the random walk started at x reaches γ before becoming negative is bounded from above by $C_\gamma(x + \sqrt{\varepsilon})$ for some constant C_γ depending on γ . As a consequence, one can construct a coupling between $Q_\gamma^\varepsilon(w)$ and $\tilde{Q}_\gamma^\varepsilon(w)$ such that

$$\mathbf{P}(Q_\gamma^\varepsilon(w) \neq \tilde{Q}_\gamma^\varepsilon(w)) \leq C_\gamma \sum_{k \in [0, 1]} |\Delta w_k|^2 (|\Delta w_k| + \sqrt{\varepsilon}). \quad (6.18)$$

Similarly, regarding the total mass of Q_γ^ε , one has the bound

$$\mathbf{E}Q_\gamma^\varepsilon(w, \mathcal{E}) = \mathbf{E}Q^\varepsilon(w, \mathcal{E}_\gamma) \leq C_\gamma \sum_{k \in [0, 1]} |\Delta w_k|^2 e^{c|\Delta w_k|}, \quad (6.19)$$

for some constant c . If w is an excursion created by the procedure above, the expected values of (6.18) and (6.19) are bounded by $C_\gamma \sqrt{\varepsilon}$ and C_γ respectively, for a possibly different constant depending on γ . Proceeding as in Lemma 5.6, it follows that, if we set

$$F_\varepsilon(w) = 1 + \frac{1}{\sqrt{\varepsilon}} \sum_{k \in [0, 1]} |\Delta w_k|^3 + \sum_{k \in [0, 1]} |\Delta w_k|^2 e^{c|\Delta w_k|},$$

we obtain the bound $\mathbf{E} \int_{\mathcal{E}} F_\varepsilon(w) \mu_\gamma^{\varepsilon, [n]}(dw) < C_{n, \gamma}$, uniformly over ε . (But this constant might potentially grow very fast as $\gamma \rightarrow 0$ and $n \rightarrow \infty$!) Combining this with (6.18), the bound (6.17) then follows at once.

As a consequence, it is sufficient to show, instead of (6.16), that

$$\lim_{\varepsilon \rightarrow 0} \mathcal{D}(\tilde{\mu}_\gamma^{\varepsilon, [n]}) = \mathcal{D}(\mu_\gamma^{[n]}). \quad (6.20)$$

This is the content of Proposition 6.20 below, which completes the proof. \square

The remainder of this section is devoted to the proof of (6.20). The outline of the proof goes as follows. First, in Section 6.2.1, we show that the law of a single excursion of the random walk (6.1), conditioned on hitting a prescribed level γ , converges as $\varepsilon \rightarrow 0$ to the Brownian excursion, conditioned to reach γ . In Section 6.2.2 we then provide a general criterion for the convergence of recursive Poisson point processes. Finally, in Section 6.2.3 we combine these results in order to obtain (6.20).

6.2.1 Convergence of excursion measures

As before, we denote by y_t the rescaled random walk given by

$$y_{(k+1)\varepsilon} = y_{k\varepsilon} + \sqrt{\varepsilon}\xi_{k+1} ,$$

where the ξ_k are an i.i.d. sequence of random variables with law ν . As before, we extend this to arbitrary times by linear interpolation. We also assume that ν has some exponential moment as before. The aim of this section is to show that if we start y_t at some initial condition $x \sim \sqrt{\varepsilon}$ and condition it on reaching a prescribed height γ before becoming negative, then its law converges to that of an unnormalised Brownian excursion, conditioned to reach height γ (call it w^γ). Throughout this whole section, we will only consider the process on a fixed time interval, which for definiteness we choose to be equal to $[0, 1]$.

A more precise description of the law \mathbf{Q}_γ of w^γ is given by the identity

$$\mathbf{Q}_\gamma(\cdot) = \frac{\int_0^\infty s^{-3/2} g_\gamma(s) \mathbf{Q}_{s,\gamma}(\cdot) ds}{\int_0^\infty s^{-3/2} g_\gamma(s) ds} ,$$

where $\mathbf{Q}_{s,\gamma}$ denotes the law of a Brownian excursion of length s , conditioned to reach level γ , and

$$g_\gamma(s) = \mathbf{Q}_{s,0}(\{w : \sup_{t \leq s} w_t \geq \gamma\}) = \mathbf{Q}(\{w : \sup_{t \leq 1} w_t \geq \gamma/\sqrt{s}\}) ,$$

where \mathbf{Q} is the standard Itô excursion measure. Since $g_\gamma(s) \rightarrow 0$ exponentially fast for small s , this does indeed define a probability measure on $\mathcal{C}(\mathbf{R}_+, \mathbf{R})$. We then turn this into a probability measure on $\mathcal{C}([0, 1], \mathbf{R})$ by restriction.

We view $\mathcal{C}([0, 1], \mathbf{R})$ as a subset of \mathcal{E} via the injection $\iota: \mathcal{C}([0, 1], \mathbf{R}) \hookrightarrow \mathcal{E}$ given by

$$\mathfrak{s}(\iota w) = 0 , \quad \mathfrak{e}(\iota w) = 1 \wedge \inf\{t > 0 : w(t) = 0\} ,$$

and by setting the path component of ιw equal to w , stopped when it reaches the time $\mathfrak{e}(\iota w)$. We endow \mathcal{E} as before with the distance d given in (5.24). Since we only deal with excursions starting at 0 and stopped before time 1, the distance d is equivalent on this set to the (pseudo-)distance \bar{d} given by

$$\bar{d}(w, w') = 1 \wedge \left(|\mathfrak{e}(w) - \mathfrak{e}(w')| + \sup_{t \in [0, 1]} |w_t - w'_t| \right) .$$

Regarding the space $\mathcal{C}([0, 1], \mathbf{R})$, we endow it throughout this section with the metric

$$d(w, w') = 1 \wedge \sup_{t \in [0, 1]} |w_t - w'_t| .$$

We furthermore denote by $\|\cdot\|_d$ the Wasserstein-1 metric associated with any distance function d , which is just the dual of the corresponding Lipschitz (semi-)norm. The main theorem of this section is given by the following:

Theorem 6.12 *Let $x > 0$ and denote by $\mathbf{Q}_{z,\gamma}^\varepsilon$ the law of the random walk (6.1), starting at $z = \sqrt{\varepsilon}x$ and conditioned to reach level γ before becoming negative. Then, for every $\bar{\beta} < \frac{1}{16}$ there exists a constant C such that*

$$\|\mathbf{Q}_{z,\gamma}^\varepsilon - \mathbf{Q}_\gamma\|_d \leq C\varepsilon^{\bar{\beta}}, \quad \|\iota^* \mathbf{Q}_{z,\gamma}^\varepsilon - \iota^* \mathbf{Q}_\gamma\|_{\bar{d}} \leq C\varepsilon^{\frac{2\bar{\beta}}{5}},$$

uniformly over all $\varepsilon \in (0, 1]$, all $x \in [0, \varepsilon^{-1/3}]$, and all $\gamma \in [\varepsilon^{1/16}, 1]$.

Our main abstract ingredient in the proof is the following criterion for the convergence of conditional probabilities when the probability of the set on which the measures are conditioned converges to 0:

Lemma 6.13 *Let μ and π be two probability measures on some Polish space \mathcal{Y} with metric d and diameter 1 and let $\mathcal{D}_\mu: \mathcal{Y} \rightarrow [0, 1]$ and $\mathcal{D}_\pi: \mathcal{Y} \rightarrow [0, 1]$. For $\varrho > 0$, set*

$$A_\varrho = \{x \in \mathcal{Y} : \exists y \in \mathcal{Y} \text{ with } \mathcal{D}_\pi(y) > \mathcal{D}_\pi(x) + d(x, y)/\varrho\},$$

$$\bar{A}_\varrho = \{x \in \mathcal{Y} : d(x, A_\varrho) \leq \varrho\}.$$

Assume that δ, ε_1 and ε_2 are such that

$$\int_{\mathcal{Y}} \mathcal{D}_\pi(x) \pi(dx) \geq \delta, \quad \|\mu - \pi\|_d \leq \varepsilon_1, \quad \sup_x |\mathcal{D}_\pi(x) - \mathcal{D}_\mu(x)| \leq \varepsilon_2,$$

where $\|\cdot\|_d$ denotes the Wasserstein-1 distance with respect to d .

Define measures $\tilde{\mu}$ and $\tilde{\pi}$ by

$$\tilde{\mu}(A) = c_\mu \int_A \mathcal{D}_\mu(x) \mu(dx), \quad \tilde{\pi}(A) = c_\pi \int_A \mathcal{D}_\pi(x) \pi(dx),$$

where c_μ and c_π are such that these are probability measures.

Then, the bound

$$\|\tilde{\mu} - \tilde{\pi}\|_d \leq \frac{1}{\delta} \left(\frac{3\varepsilon_1}{\varrho} + \varepsilon_2 + 2\pi(\bar{A}_\varrho) \right), \quad (6.21)$$

holds for every $\varrho \leq 1$. In particular, one has $\int_{\mathcal{Y}} \mathcal{D}_\mu(x) \mu(dx) > 0$ whenever the right hand side in (6.21) is strictly smaller than 1, so that the bound is non-trivial.

Proof. Let $f: \mathcal{Y} \rightarrow \mathbf{R}$ be a test function such that $\text{Lip}_d(f) \leq 1$. Since the diameter of \mathcal{Y} is assumed to be 1, we can assume without loss of generality (by possibly adding a constant to f) that $\sup_x |f(x)| \leq \frac{1}{2}$. Since one has the identity

$$\|\tilde{\mu} - \tilde{\pi}\|_d = \sup_{\text{Lip}_d(f) \leq 1} \left(c_\mu \int f(x) \mathcal{D}_\mu(x) \mu(dx) - c_\pi \int f(x) \mathcal{D}_\pi(x) \pi(dx) \right) = \sup_{\text{Lip}_d(f) \leq 1} \mathcal{I}_f,$$

our aim is to bound \mathcal{I}_f , uniformly over f . Note first that, by the bound on f ,

$$\begin{aligned} \mathcal{I}_f &\leq \frac{c_\pi}{2} \left| \frac{1}{c_\pi} - \frac{1}{c_\mu} \right| + \frac{c_\pi}{2} \int |\mathcal{D}_\mu(x) - \mathcal{D}_\pi(x)| \mu(dx) \\ &\quad + c_\pi \left| \int f(x) \mathcal{D}_\pi(x) \mu(dx) - \int f(x) \mathcal{D}_\pi(x) \pi(dx) \right|. \end{aligned}$$

Note first that the first term is nothing but a particular instance of the last term with $f = \frac{1}{2}$. Since the second term is furthermore trivially bounded by $\varepsilon_2/(2\gamma)$, it suffices to bound the last term.

The problem in bounding this term arises of course from the fact that \mathcal{D}_π is not Lipschitz continuous. For any $\varrho > 0$, we can however “mollify” this function by setting

$$\mathcal{D}_\pi^\varrho(x) = \sup_{y \in \mathcal{Y}} \left(\mathcal{D}_\pi(y) - \frac{d(x, y)}{\varrho} \right).$$

It is then straightforward to check that $\text{Lip}_d(\mathcal{D}_\pi^\varrho) \leq \varrho^{-1}$, that $\mathcal{D}_\pi(x) \leq \mathcal{D}_\pi^\varrho(x) \leq \sup_y \mathcal{D}_\pi(y)$, and that furthermore $\mathcal{D}_\pi^\varrho(x) = \mathcal{D}_\pi(x)$ for all $x \notin A_\varrho$. It then follows from the definition of ε_1 that

$$\left| \int f(x) \mathcal{D}_\pi^\varrho(x) \mu(dx) - \int f(x) \mathcal{D}_\pi^\varrho(x) \pi(dx) \right| \leq \varepsilon_1 (1 + (2\varrho)^{-1}).$$

Furthermore,

$$\left| \int f(x) \mathcal{D}_\pi^\varrho(x) \pi(dx) - \int f(x) \mathcal{D}_\pi(x) \pi(dx) \right| \leq \frac{\pi(A_\varrho)}{2},$$

and similarly for the term with π replaced by μ . In order to bound $\mu(A_\varrho)$, we set as above

$$F^\varrho(x) = \sup_{y \in \mathcal{Y}} \left(\mathbf{1}_{A_\varrho}(y) - \frac{d(x, y)}{\varrho} \right),$$

so that

$$\mu(A_\varrho) \leq \int F^\varrho(x) \mu(dx) \leq \frac{\varepsilon_1}{\varrho} + \int F^\varrho(x) \pi(dx) \leq \frac{\varepsilon_1}{\varrho} + \pi(\bar{A}_\varrho),$$

where we used the fact that F^ϱ vanishes outside of \bar{A}_ϱ . Collecting all of these bounds, the claim follows. \square

An alternative description of w^γ is given by the following. Let Y be a Bessel-3 process starting at the origin and let τ_γ be its first hitting time of γ , i.e. $\tau_\gamma = \inf\{t \geq 0 : Y_t \geq \gamma\}$. Let furthermore B be a Brownian motion independent of Y , which is stopped when it reaches the level γ . Then, one has the decomposition

$$w_t^\gamma = \begin{cases} Y_t & \text{for } t \leq \tau_\gamma, \\ \gamma - B_{t-\tau_\gamma} & \text{for } t \geq \tau_\gamma. \end{cases} \quad (6.22)$$

This is a consequence of the symmetry of the Brownian excursion under time reversal, combined with [RW94, Theorem 49.1] for example.

We can use the above decomposition to obtain the following bound:

Lemma 6.14 *For any $\beta < \frac{1}{4}$, there exists a constant C such that, for every $\gamma, \gamma' \in (0, 1]$, one has the bound*

$$\|\mathbf{Q}_\gamma - \mathbf{Q}_{\gamma'}\|_d \leq C |\gamma - \gamma'|^\beta.$$

Proof. The decomposition (6.22) suggests a natural coupling between w^γ and $w_{\gamma'}$ by building them from the same basic building blocks Y and B . The characterisation of

the Bessel-3 process as the norm of a 3-dimensional Brownian motion, together with standard hitting estimates for Brownian motion, imply that

$$\mathbf{P}(|\tau_\gamma - \tau_{\gamma'}| > \zeta) \leq 1 \wedge \frac{C|\gamma - \gamma'|}{\zeta^{3/2}},$$

so that in particular $\mathbf{P}(|\tau_\gamma - \tau_{\gamma'}| \geq \sqrt{|\gamma - \gamma'|}) \leq C|\gamma - \gamma'|^{1/4}$. The result now follows from the fact that both B and Y are almost surely α -Hölder continuous for every $\alpha < \frac{1}{2}$. \square

Lemma 6.15 *Let B_z^γ be a Brownian motion starting at z , conditioned to hit γ before 0, and stopped upon its return to 0. Then, for every $\beta < 1$, there exists a constant C depending on β such that the bound*

$$\|\mathcal{D}(B_\varepsilon^\gamma) - \mathbf{Q}_\gamma\|_d \leq C_\gamma \varepsilon^\beta,$$

holds uniformly over $\varepsilon \in (0, \gamma \wedge 1]$ and $\gamma > 0$.

Proof. Let w^γ be as above and let τ_ε be the first passage time of w^γ through ε . Then, it follows from the decomposition (6.22) that one has the exact identity

$$B_\varepsilon^\gamma(\cdot) \stackrel{\text{law}}{=} w^\gamma(\cdot - \tau_\varepsilon). \quad (6.23)$$

It then follows from the small ball estimates of Brownian motion that, for every $\bar{\zeta} < 2$, one has the bound

$$\mathbf{P}(\tau_\varepsilon \geq \varepsilon^{\bar{\zeta}}) \lesssim \varepsilon.$$

The desired estimate then follows at once from the Hölder regularity of w^γ . \square

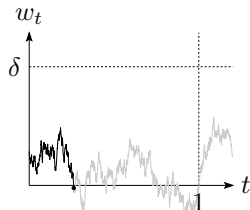
Proposition 6.16 *Let $\alpha \in (0, \frac{1}{8})$. Suppose further that $\gamma > 0$ is fixed and denote by y_t^γ the random walk y_t conditioned to hit $[\gamma, \infty)$ before hitting \mathbf{R}_- and stopped when it then hits \mathbf{R}_- . Assume that $y_0^\gamma = \varepsilon^\alpha$. Then the law of y^γ converges weakly to \mathbf{Q}_γ as $\varepsilon \rightarrow 0$. Furthermore, for every $\beta > 0$, there exists a constant C such that the bound*

$$\|\mathcal{D}(y^\gamma) - \mathbf{Q}_\gamma\|_d \leq C(\sqrt{\gamma} \varepsilon^{\frac{1}{8}-\alpha} + \varepsilon^{\alpha-\beta}),$$

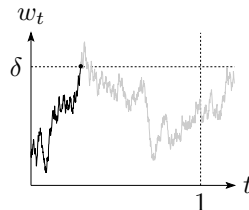
holds uniformly over all $\varepsilon \leq 1$ and $\gamma \in [\varepsilon^\alpha, 1]$.

Proof. By Lemma 6.15, it suffices to compare the law of y^γ with that of $B_{\varepsilon^\alpha}^\gamma$, which is itself going to be a consequence of Lemma 6.13.

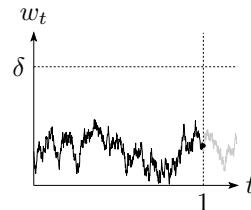
To see this, we partition the state space $\mathcal{Y} = \{w \in \mathcal{C}([0, 1], \mathbf{R}) : w_0 = \varepsilon^\alpha\}$ into three sets in the following way. Given a continuous function w with $w_0 \in (0, \gamma)$, we set $\tau = 1 \wedge \inf\{t > 0 : w_t \notin [0, \gamma]\}$, and we define sets $A^{(i)}$ with $i \in \{1, 2, 3\}$ by



$A^{(1)} = \{w : w_\tau = 0\},$



$A^{(2)} = \{w : w_\tau = \gamma\},$



$A^{(3)} = \{w : \tau = 1\}.$

Define furthermore functions F_γ and F_γ^ε on \mathcal{Y} by

$$F_\gamma(w) = \frac{w_\tau}{\gamma}, \quad F_\gamma^\varepsilon(w) = \begin{cases} 0 & \text{if } w \in A^{(1)}, \\ 1 & \text{if } w \in A^{(2)}, \\ \bar{P}_{\frac{\gamma}{\sqrt{\varepsilon}}, \frac{w_1}{\sqrt{\varepsilon}}} & \text{if } w \in A^{(3)}, \end{cases}$$

where $\bar{P}_{z,\gamma}$ is defined as in (5.13). With these definitions at hand, if we set $\mu = \mathcal{D}(y)$ with $y_0 = \varepsilon^\alpha$, $\pi = \mathcal{D}(B_{\varepsilon^\alpha})$, $\mathcal{D}_\mu = F_\gamma^\varepsilon$, and $\mathcal{D}_\pi = F_\gamma$, then we are precisely in the setting of Lemma 6.13 with $\tilde{\mu} = \mathcal{D}(y^\gamma)$ and $\tilde{\pi} = \mathcal{D}(B_{\varepsilon^\alpha}^\gamma)$.

Note first that, since F_γ is precisely the probability that a Brownian motion started from w_1 hits γ before 0, we have

$$\delta = \int \mathcal{D}_\pi(w) \pi(dw) = \frac{\varepsilon^\alpha}{\gamma}.$$

Furthermore, it follows from Corollary 5.2 (and from the fact that F_γ and F_γ^ε coincide outside of $A^{(3)}$) that

$$\varepsilon_1 = \sup_w |F_\gamma(w) - F_\gamma^\varepsilon(w)| \lesssim \frac{\varepsilon^{\frac{1}{4}}}{\gamma}.$$

(We could have replaced $\frac{1}{4}$ by any exponent less than $\frac{1}{2}$ here.) Regarding the distance between the unconditioned measures, we obtain from [Fra73] the bound

$$\varepsilon_2 = \|\mu - \pi\|_d \lesssim \varepsilon^{\frac{1}{8}}.$$

It thus remains to obtain a bound on \bar{A}_ϱ . By the definition of A_ϱ and of F_γ , $w \in A_\varrho$ implies that either $w \in A^{(1)} \cup A^{(3)}$ and $d(w, A^{(2)}) \leq \varrho$ or $w \in A^{(1)}$ and $d(w, A^{(3)}) \leq \varrho$. This implies that

$$A_\varrho \subset \left\{ w : \sup_{t \in [0,1]} w_t \in [\gamma - \varrho, \gamma] \right\} \cup \left\{ w : \inf_{t \in [0,1]} w_t \in [-\varrho, 0] \right\},$$

so that

$$\bar{A}_\varrho \subset \left\{ w : \sup_{t \in [0,1]} w_t \in [\gamma - 2\varrho, \gamma + \varrho] \right\} \cup \left\{ w : \inf_{t \in [0,1]} w_t \in [-2\varrho, \varrho] \right\}.$$

Since (by the reflection principle) the law of the extremum of Brownian motion over a finite time interval has a smooth density with respect to Lebesgue measure, there exists a constant C independent of ε and δ such that $\pi(\bar{A}_\varrho) \leq C\varrho$.

Inserting these bounds into Lemma 6.13, we thus obtain the bound

$$\|\mathcal{D}(y^\gamma) - \mathcal{D}(B_{\varepsilon^\alpha}^\gamma)\|_d \lesssim \frac{\gamma}{\varepsilon^\alpha} \left(\frac{\varepsilon^{\frac{1}{4}}}{\varrho\gamma} + \varepsilon^{\frac{1}{8}} + \varrho \right).$$

Setting $\varrho = \varepsilon^{\frac{1}{8}}\gamma^{-\frac{1}{2}}$ completes the proof. \square

We now have all the ingredients required for the proof of Theorem 6.12.

Proof of Theorem 6.12. Assume as in previous proof that y_t^γ is the random walk y_t conditioned to hit $[\gamma, \infty)$ before hitting \mathbf{R}_- and stopped when it then hits \mathbf{R}_- . In contrast to the above setup we now assume that $y_0^\gamma = z\sqrt{\varepsilon}$ for some $z \geq 0$. Let k_0 be given by

$$k_0 = \inf \{k > 0 : y_{k\varepsilon}^\gamma \geq \varepsilon^{\frac{1}{16}}\}.$$

It then follows from [Fra73], combined with standard small ball estimates for Brownian motion that, for every $\beta < \frac{1}{8}$ and every $p > 0$ there exists a constant C such that

$$\mathbf{P}(\varepsilon k_0 > \varepsilon^\beta) \lesssim \varepsilon^p, \quad (6.24)$$

uniformly over $\varepsilon \leq 1$, for all z such that $z\sqrt{\varepsilon} \leq \varepsilon^{\frac{1}{16}}$. Furthermore, the probability that $y_{k_0\varepsilon}^\gamma > 2\varepsilon^{\frac{1}{16}}$ (say) is exponentially small in ε , again uniformly over x . It then follows from Proposition 6.16 (choosing $\alpha = \frac{1}{16}$) that, for every $\bar{\beta} < \frac{1}{16}$ one can construct a joint realisation of y^γ and w^γ such that

$$\mathbf{E}d(y^\gamma, w^\gamma(\cdot - \varepsilon k_0)) \lesssim \varepsilon^{\bar{\beta}}.$$

(Here we extend w^γ by setting it to 0 for negative times.) On the other hand, the Hölder regularity of the sample paths of w^γ together with (6.24) implies that

$$\mathbf{E}d(w^\gamma, w^\gamma(\cdot - \varepsilon k_0)) \lesssim \varepsilon^{\bar{\beta}},$$

so that the bound on $\|\mathbf{Q}_{z,\gamma}^\varepsilon - \mathbf{Q}_\gamma\|_d$ follows.

In order to obtain the bound on $\|\iota^* \mathbf{Q}_{z,\gamma}^\varepsilon - \iota^* \mathbf{Q}_\gamma\|_{\bar{d}}$, we make use of the same coupling between y^γ and w^γ as above, so that $\mathbf{E}d(y^\gamma, w^\gamma) \lesssim \varepsilon^{\bar{\beta}}$. It then remains to obtain a bound on

$$|\mathbf{e}(\iota y^\gamma) - \mathbf{e}(\iota w^\gamma)|.$$

For this, note first that, by Chebychev's inequality, one has the bound

$$\mathbf{P}(d(y^\gamma, w^\gamma) \geq \varepsilon^\beta) \lesssim \varepsilon^{\bar{\beta}-\beta}, \quad (6.25)$$

valid for every $\beta \leq \bar{\beta}$. Consider now any two paths y^γ and w^γ at distance less than ε^β and define

$$\tau_1 = 1 \wedge \inf\{t > 0 : w^\gamma(t) \leq \varepsilon^\beta\}, \quad \tau_2 = 1 \wedge \inf\{t > \tau_1 : w^\gamma(t) < -\varepsilon^\beta\}.$$

In this way, one has both $\mathbf{e}(\iota w^\gamma) \in [\tau_1, \tau_2]$ and $\mathbf{e}(\iota y^\gamma) \in [\tau_1, \tau_2]$, so that it remains to obtain a bound on $\tau_2 - \tau_1$. The explicit expression for the hitting time of a line for a Brownian motion yields

$$\mathbf{P}(\tau_2 - \tau_1 \geq \varepsilon^\alpha) \leq \varepsilon^{\beta-\frac{\alpha}{2}},$$

for any $\alpha < 2\beta$. Choosing $\beta = \frac{3\bar{\beta}}{5}$ and $\alpha = \frac{2\bar{\beta}}{5}$ and combining this with (6.25), we then obtain

$$\mathbf{P}(|\mathbf{e}(\iota y^\gamma) - \mathbf{e}(\iota w^\gamma)| \geq \varepsilon^{\frac{2\bar{\beta}}{5}}) \lesssim \varepsilon^{\frac{2\bar{\beta}}{5}},$$

from which the bound follows. \square

6.2.2 Convergence of recursive Poisson point processes

The aim of this section is to provide a general result allowing us to bound the distance between two recursive Poisson point processes of the same depth n in terms of their respective kernels. This result is the main abstract result on which the proof of the convergence result (6.20) will be based. One difficulty that we have to overcome is that there is very little uniformity in the proximity of the kernel describing $\mu_\gamma^{\varepsilon,[n]}$ to the one describing $\mu_\gamma^{[n]}$.

Throughout this section, given a Polish space \mathcal{X} with a distance function d bounded by 1, we define the Wasserstein-1 distance between any two positive measures with finite mass (and not just probability measures!) by

$$\|\mu - \pi\|_1 = \sup_{\substack{\text{Lip } f \leq 1 \\ \|f\|_\infty \leq 1}} \left(\int_{\mathcal{X}} f(x) \mu(dx) - \int_{\mathcal{X}} f(x) \pi(dx) \right).$$

If μ and π happen to have the same mass, then the additional constraint on the supremum norm $\|f\|_\infty$ of f is redundant in the above expression, and we simply recover the usual Wasserstein-1 distance. In the case where the masses of μ and π are different however, our choice of definitions ensures that

$$\|\mu - \pi\|_1 \approx |\mu(\mathcal{X}) - \pi(\mathcal{X})| + |\mu(\mathcal{X})| \left\| \frac{\mu}{\mu(\mathcal{X})} - \frac{\pi}{\pi(\mathcal{X})} \right\|_1, \quad (6.26)$$

where \approx denotes that both quantities are bounded by multiples of each other, with proportionality constants that are independent of μ and π .

The main result of this section is the following:

Theorem 6.17 *Let \mathcal{Q} and $\bar{\mathcal{Q}}$: $\mathcal{X} \rightarrow \mathbb{M}_+(\mathcal{X})$ be two measurable maps and assume that the Polish space \mathcal{X} is endowed with a metric d bounded by 1. Let $A \subset \mathcal{X}$, $\varepsilon \in (0, 1]$ and $K \geq 1$ be such that the bounds*

$$\sup_{x \in \mathcal{X}} \mathcal{Q}(x, \mathcal{X}) \leq K, \quad \|\mathcal{Q}(x, \mathcal{X}) - \mathcal{Q}(y, \mathcal{X})\|_1 \leq Kd(x, y), \quad (6.28)$$

$$\sup_{x \in A} \bar{\mathcal{Q}}(x, A^c) \leq \varepsilon, \quad \sup_{x \in A} \|\mathcal{Q}(x) - \bar{\mathcal{Q}}(x)\|_1 \leq \varepsilon, \quad (6.29)$$

hold, where we use the notation $A^c = \mathcal{X} \setminus A$.

Fix $n > 0$, $\bar{x} \in A$ and $x \in \mathcal{X}$, and denote by $\mu_x^{[n]}$ and $\bar{\mu}_{\bar{x}}^{[n]}$ the recursive Poisson point processes with respective kernels \mathcal{Q} and $\bar{\mathcal{Q}}$. Then, there exists a coupling between $\mu_x^{[n]}$ and $\bar{\mu}_{\bar{x}}^{[n]}$ such that

$$\mathbf{E}(1 \wedge \|\mu_x^{[n]} - \bar{\mu}_{\bar{x}}^{[n]}\|_1) \lesssim C(\sqrt{\varepsilon} + d(x, \bar{x})),$$

where the proportionality constant C depends only on K and n .

Remark 6.18 One useful feature of the way that we set up the bounds in the statement is that we only require information about the kernel $\bar{\mathcal{Q}}$ on the set A . In the application we have in mind, the kernel $\bar{\mathcal{Q}}$ will be the one describing $\tilde{\mu}_\gamma^{\varepsilon, [n]}$, while the set A will consist of trajectories exhibiting “typical” behaviour at small scales.

Before we turn to the proof of this theorem, we show that if $\pi_n \rightarrow \pi$ in the Wasserstein-1 sense, then the (usual) Poisson point processes with these intensity measures also converge to each other weakly in the Wasserstein-1 distance:

Proposition 6.19 *Let π and $\bar{\pi}$ be two finite positive measures on a Polish space \mathcal{X} endowed with a metric $d \leq 1$ and let μ and $\bar{\mu}$ be the corresponding Poisson point processes on \mathcal{X} . Then, there exists a constant C and a coupling between μ and $\bar{\mu}$ such that*

$$\mathbf{E}(1 \wedge \|\mu - \bar{\mu}\|_1) \leq C(\|\pi - \bar{\pi}\|_1 \wedge 1).$$

Proof. The proof relies on the fact that, if $\mathcal{P}(\lambda)$ denotes the Poisson distribution with parameter λ , one has the total variation bound

$$d_{\text{TV}}(\mathcal{P}(\lambda), \mathcal{P}(\bar{\lambda})) \leq |\lambda - \bar{\lambda}| \wedge 1, \quad (6.30)$$

see for example [AJ06].

We can construct μ (and similarly for $\bar{\mu}$) in the following way. First, draw a Poisson random variable N with parameter $\pi(\mathcal{X})$. Then, draw N independent random variables $\{X_1, \dots, X_N\}$ with law $\pi/\pi(\mathcal{X})$ and set

$$\mu = \sum_{k=1}^N \delta_{X_k}. \quad (6.31)$$

By (6.30), we can now construct a Poisson random variable \bar{N} with parameter $\bar{\pi}(\mathcal{X})$ such that

$$\mathbf{P}(\bar{N} \neq N) \leq |\pi(\mathcal{X}) - \bar{\pi}(\mathcal{X})|. \quad (6.32)$$

Assuming that $\bar{N} = N$, we can then draw random variables $\{\bar{X}_1, \dots, \bar{X}_N\}$ in such a way that the pairs (\bar{X}_k, X_k) are distributed according to a coupling between $\pi/\pi(\mathcal{X})$ and $\bar{\pi}/\bar{\pi}(\mathcal{X})$ that minimises their expected distance. If $\bar{N} \neq N$, then we simply draw $\{\bar{X}_1, \dots, \bar{X}_N\}$ according to $\bar{\pi}/\bar{\pi}(\mathcal{X})$, independently of the X_k .

If we then define $\bar{\mu}$ similarly to (6.31), it follows that

$$\|\mu - \bar{\mu}\|_1 \leq \begin{cases} N \left\| \frac{\pi}{\pi(\mathcal{X})} - \frac{\bar{\pi}}{\bar{\pi}(\mathcal{X})} \right\|_1 & \text{if } N = \bar{N}, \\ N + \bar{N} & \text{otherwise.} \end{cases}$$

As a consequence, we obtain the bound

$$\mathbf{E}(1 \wedge \|\mu - \bar{\mu}\|_1) \leq \mathbf{P}(N \neq \bar{N}) + \left\| \frac{\pi}{\pi(\mathcal{X})} - \frac{\bar{\pi}}{\bar{\pi}(\mathcal{X})} \right\|_1 \mathbf{E}N,$$

so that the claim follows from (6.32), the definition of N , and (6.26). \square

Proof of Theorem 6.17. Note first that, by combining the first bound in (6.28) with the second bound in (6.29), we obtain the bound

$$\sup_{x \in A} \bar{Q}(x, \mathcal{X}) \leq K + \varepsilon.$$

It follows that we have the recursive bound

$$\mathbf{E}(\bar{\mu}_x^n(A) \mid \bar{\mu}_x^{[n-1]}(A^c) = 0) \leq 2K \mathbf{E}(\bar{\mu}_x^{n-1}(A) \mid \bar{\mu}_x^{[n-2]}(A^c) = 0),$$

so that, since $\bar{\mu}_x^0(A) = 1$ by assumption,

$$\mathbf{E}(\bar{\mu}_x^n(A) \mid \bar{\mu}_x^{[n-1]}(A^c) = 0) \leq (2K)^n. \quad (6.33)$$

On the other hand, defining the positive measures π^n and $\bar{\pi}^n$ by

$$\pi^{n+1} = \int_{\mathcal{X}} \bar{Q}(y) \mu_x^n(dy) \quad \text{and} \quad \bar{\pi}^{n+1} = \int_{\mathcal{X}} \bar{Q}(y) \bar{\mu}_x^n(dy)$$

we have the inequality

$$\begin{aligned}
\mathbf{P}(\bar{\mu}_x^{[n]}(A^c) > 0) &\leq \mathbf{P}(\bar{\mu}_x^{[n-1]}(A^c) > 0) + \mathbf{P}(\bar{\mu}_x^n(A^c) > 0 \mid \bar{\mu}_x^{[n-1]}(A^c) = 0) \\
&\leq \mathbf{P}(\bar{\mu}_x^{[n-1]}(A^c) > 0) + \mathbf{E}(\bar{\pi}^n(A^c) \mid \bar{\mu}_x^{[n-1]}(A^c) = 0) \\
&\leq \mathbf{P}(\bar{\mu}_x^{[n-1]}(A^c) > 0) + \varepsilon D + \mathbf{P}(\bar{\mu}_x^{n-1}(A) > D \mid \bar{\mu}_x^{[n-1]}(A^c) = 0) \\
&\leq \mathbf{P}(\bar{\mu}_x^{[n-1]}(A^c) > 0) + \varepsilon D + \frac{1}{D} \mathbf{E}(\bar{\mu}_x^{n-1}(A) \mid \bar{\mu}_x^{[n-1]}(A^c) = 0) \\
&\leq \mathbf{P}(\bar{\mu}_x^{[n-1]}(A^c) > 0) + \varepsilon D + \frac{(2K)^{n-1}}{D},
\end{aligned}$$

which is valid uniformly over all $D > 0$. Choosing $D \sim 1/\sqrt{\varepsilon}$, we thus obtain the recursion relation

$$\mathbf{P}(\bar{\mu}_x^{[n]}(A^c) > 0) \leq \mathbf{P}(\bar{\mu}_x^{[n-1]}(A^c) > 0) + C\sqrt{\varepsilon},$$

from which it follows that

$$\mathbf{P}(\bar{\mu}_x^{[n]}(A^c) > 0) \leq C\sqrt{\varepsilon}, \quad (6.34)$$

where in both cases the constant C depends on K and n , but not on ε .

Note now that we have the bound

$$\begin{aligned}
\|\pi^{n+1} - \bar{\pi}^{n+1}\|_1 &\leq \left\| \int_{\mathcal{X}} \mathcal{Q}(y) (\mu_x^n - \bar{\mu}_x^n)(dy) \right\|_1 + \int_{\mathcal{X}} \|\mathcal{Q}(y) - \bar{\mathcal{Q}}(y)\|_1 \bar{\mu}_x^n(dy) \\
&\leq \|\mu_x^n - \bar{\mu}_x^n\|_1 (\text{Lip}(\mathcal{Q}) + \|\mathcal{Q}\|_\infty) + \varepsilon \bar{\mu}_x^n(A) \\
&\quad + \int_{A^c} \|\mathcal{Q}(y) - \bar{\mathcal{Q}}(y)\|_1 \bar{\mu}_x^n(dx),
\end{aligned}$$

so that the bound

$$1 \wedge \|\pi^{n+1} - \bar{\pi}^{n+1}\|_1 \leq 2K(\|\mu_x^n - \bar{\mu}_x^n\|_1 \wedge 1) + \varepsilon D + \mathbf{1}_{\bar{\mu}_x^n(A) > D} + \mathbf{1}_{\bar{\mu}_x^n(A^c) > 0},$$

is valid for every $D > 0$. Furthermore, by Chebyshev's inequality and (6.33)–(6.34), one has

$$\begin{aligned}
\mathbf{P}(\bar{\mu}_x^n(A) > D) &\leq \mathbf{P}(\bar{\mu}_x^n(A) > D \mid \bar{\mu}_x^{n-1}(A^c) = 0) + \mathbf{P}(\bar{\mu}_x^{n-1}(A^c) > 0) \\
&\leq \frac{K^n}{D} + C\sqrt{\varepsilon}.
\end{aligned}$$

Choosing again $D \sim \sqrt{\varepsilon}$, we thus obtain the bound

$$\mathbf{E}(1 \wedge \|\pi^{n+1} - \bar{\pi}^{n+1}\|_1) \leq 2K\mathbf{E}(\|\mu_x^n - \bar{\mu}_x^n\|_1 \wedge 1) + C\sqrt{\varepsilon}.$$

Applying Proposition 6.19, we conclude that, given μ_x^n and $\bar{\mu}_x^n$, it is possible to construct a coupling between μ_x^{n+1} and $\bar{\mu}_x^{n+1}$ such that

$$\mathbf{E}(1 \wedge \|\mu_x^{n+1} - \bar{\mu}_x^{n+1}\|_1) \lesssim \mathbf{E}(1 \wedge \|\mu_x^n - \bar{\mu}_x^n\|_1) + \sqrt{\varepsilon},$$

from whence the claim now follows at once. \square

6.2.3 Convergence of the truncated distributions

We are now in a position to provide the proof of (6.20). Again, throughout this section, we make the standing assumption that the one-step distribution ν for the random walk (6.1) has some exponential moment. We also use the notations $\tilde{\mu}_\gamma^{\varepsilon, [n]}$ and $\mu_\gamma^{[n]}$ from (6.20). Let \mathcal{E} be the space of real-valued excursions as before. We then introduce the map $J^\varepsilon : \mathcal{E} \rightarrow M_+(\mathbf{R}^2)$ given by

$$J^\varepsilon(w)(dz, dt) = a\varepsilon \sum_{\varepsilon k \in [0, 1]} e^{a\sqrt{\varepsilon}z} \delta_{\varepsilon k}(dt) \mathbf{1}_{[0, \Delta_k^\varepsilon w]}(z) dz, \quad (6.35)$$

where δ_z denotes the Dirac measure located at z , $\Delta_k^\varepsilon w$ is defined by

$$\Delta_k^\varepsilon w = \frac{w_{(k+1)\varepsilon} - w_{k\varepsilon}}{\sqrt{\varepsilon}},$$

and we used the convention that $\mathbf{1}_{[0, z]} = 0$ if $z < 0$.

As before, denote by $\mathbf{Q}_{z, \gamma}^\varepsilon$ the law of the random walk (6.1), starting at $z\sqrt{\varepsilon}$, and conditioned to hit level γ before becoming negative. We stop it as soon as it hits \mathbf{R}_- , so that we interpret $\mathbf{Q}_{z, \gamma}^\varepsilon$ as a measure on \mathcal{E}_0 . Recall also that $\bar{P}_{z, \gamma/\sqrt{\varepsilon}}$, with $\bar{P}_{z, \gamma}$ defined as in Section 5.1.2 denotes the probability that this event actually happens.

With this notation, the measure $\mathcal{Q}_\gamma^\varepsilon$ describing $\tilde{\mu}_\gamma^{\varepsilon, [n]}$ (i.e. $\mathcal{Q}_\gamma^\varepsilon(w, \cdot)$ is the intensity measure of $\tilde{Q}_\gamma^\varepsilon(w)$) is given by

$$\begin{aligned} \mathcal{Q}_\gamma^\varepsilon(w, \cdot) &= \int_{\mathbf{R}^2} \Theta_{w, t}^* \mathbf{Q}_{z, \gamma}^\varepsilon \bar{P}_{z, \gamma/\sqrt{\varepsilon}} \frac{J^\varepsilon(w)(dz, dt)}{\sqrt{\varepsilon}} \\ &= \int_{\mathbf{R}^2} \left(\frac{1}{\gamma} \Theta_{w, t}^* \mathbf{Q}_{z, \gamma}^\varepsilon \right) \left(\frac{\gamma}{\sqrt{\varepsilon}} \bar{P}_{z, \gamma/\sqrt{\varepsilon}} \right) J^\varepsilon(w)(dz, dt). \end{aligned}$$

Note now that if w is a typical realisation of $\mathbf{Q}_{z, \gamma}^\varepsilon$, then $J^\varepsilon(w)$ is expected to be close to the measure

$$J(w) = \mathbf{1}_{\{(w) \cap [0, 1]\}}(t) dt \otimes \hat{\nu}(dz),$$

where $\hat{\nu}$ is the measure on \mathbf{R}_+ given by

$$\int G(z) \hat{\nu}(dz) = \int_0^\infty \int_0^z G(y) dy \nu(dz),$$

for any test function G . This is because $e^{a\sqrt{\varepsilon}\Delta_k^\varepsilon w} \sim 1$ and the law of $\Delta_k^\varepsilon w$ would be given by ν , were it not for the conditioning.

On the other hand, the kernel \mathcal{Q}_γ describing the truncated Brownian fan $\mu_\gamma^{[n]}$ with parameter a is given by

$$\mathcal{Q}_\gamma(w, \cdot) = \frac{a}{2\delta} \int_{\{(w) \cap [0, 1]\}} \Theta_{w, t}^* \mathbf{Q}_\gamma dt = \frac{1}{\delta} \int_{\mathbf{R}^2} \Theta_{w, t}^* \mathbf{Q}_\gamma G(z) J(w)(dz, dt), \quad (6.36)$$

where \mathbf{Q}_γ is the law of a Brownian excursion conditioned to reach level γ . This is the case because of the well-known fact that $\frac{1}{\gamma} \mathbf{Q}_\gamma$ is the law of the unnormalised Brownian excursion *restricted* to the set of excursions reaching level γ . The second identity in (6.36) is a consequence of the definition of $\hat{\nu}$, combined with Proposition 5.3. With these notations at hand, we have:

Proposition 6.20 *Consider the setting and assumptions of Theorem 6.11. For every $\delta < \frac{1}{32}$, there exists a constant C depending on γ , n and δ such that*

$$\mathbf{E}(1 \wedge \|\tilde{\mu}_\gamma^{\varepsilon, [n]} - \mu_\gamma^{[n]}\|_1) \leq C\varepsilon^\delta,$$

uniformly over $\varepsilon \leq \varepsilon_0$ for some ε_0 small enough.

Proof. We apply Theorem 6.17 with $\mathcal{Q} = \mathcal{Q}_\gamma$, $\bar{\mathcal{Q}} = \mathcal{Q}_\gamma^\varepsilon$, and A to be determined. Using the results obtained earlier in this section, it turns out that the assumptions are then relatively straightforward to check.

For convenience, we introduce the notation

$$\tilde{J}^\varepsilon(w)(dz, dt) = G(x)\tilde{J}^\varepsilon(w)(dz, dt),$$

and similarly for \tilde{J} . We also denote by $\Pi_2: \mathbf{R}^2 \rightarrow \mathbf{R}$ the projection onto the second component, so that $\Pi_2^* \tilde{J}^\varepsilon(w)$ (and similarly for \tilde{J}) is the projection of $\tilde{J}^\varepsilon(w)$ onto the t -component. We also denote by $\Omega_\varepsilon \subset \mathcal{E}$ the subset of excursions given by

$$\Omega_\varepsilon = \Omega_\varepsilon^{(1)} \cap \Omega_\varepsilon^{(2)},$$

where we set

$$\begin{aligned} \Omega_\varepsilon^{(1)} &= \{w \in \mathcal{E} : \tilde{J}^\varepsilon(w)(\{x > \varepsilon^{-1/3}\}) = 0\}, \\ \Omega_\varepsilon^{(2)} &= \{w \in \mathcal{E} : \|\Pi_2^* \tilde{J}(w) - \Pi_2^* \tilde{J}^\varepsilon(w)\|_1 \leq \varepsilon^{1/10}\}. \end{aligned}$$

Also, in the definition of $\Omega_\varepsilon^{(2)}$, the Wasserstein-1 distance is taken with respect to the somewhat unusual distance on \mathbf{R} given by

$$d(t, t') = 1 \wedge |t - t'|^{1/3}. \quad (6.37)$$

The reason for this seemingly strange choice will become clear later. The set Ω_ε defined in this way will play the role of A when applying Theorem 6.17.

Note first that $\mathcal{Q}_\gamma(w, \mathcal{X}) \leq \frac{a}{2\gamma}$, which is bounded independently of w . Furthermore, one has the bound

$$\begin{aligned} \|\mathcal{Q}_\gamma(w, \mathcal{X}) - \mathcal{Q}_\gamma(\bar{w}, \mathcal{X})\|_1 &\leq \frac{a}{2\gamma} \int_{\mathfrak{l}(w) \cap \mathfrak{l}(\bar{w}) \cap [0, 1]} \|\Theta_{w, t}^* \mathbf{Q}_\gamma - \Theta_{\bar{w}, t}^* \mathbf{Q}_\gamma\|_1 dt \\ &\quad + \frac{a}{2\gamma} (|\mathfrak{s}(w) - \mathfrak{s}(\bar{w})| + |\mathfrak{e}(w) - \mathfrak{e}(\bar{w})|). \end{aligned}$$

Since furthermore $\|\Theta_{w, t}^* \mathbf{Q}_\gamma - \Theta_{\bar{w}, t}^* \mathbf{Q}_\gamma\|_1 \leq |w_t - \bar{w}_t|$, it does indeed follow that \mathcal{Q} is globally Lipschitz continuous as required, so that (6.28) holds with some constant K depending on γ .

It remains to check that (6.29) holds, with ε replaced by some power of ε . By Theorem 6.12, we already know that there exists a constant C , possibly depending on γ , such that

$$\|\mathbf{Q}_{z, \gamma}^\varepsilon - \mathbf{Q}_\gamma\|_1 \leq C\varepsilon^\delta,$$

for any exponent $\delta \in (0, \frac{1}{16})$, and uniformly over all $z \leq \varepsilon^{-1/3}$. As a consequence, for every $w \in \Omega_\varepsilon$, we have the bound

$$\|\mathcal{Q}_\gamma(w, \mathcal{X}) - \mathcal{Q}_\gamma^\varepsilon(w, \mathcal{X})\|_1 \leq \frac{1}{\gamma} \int_{\mathbf{R}^2} \left(\frac{\gamma}{\sqrt{\varepsilon}} \bar{P}_{z, \frac{\gamma}{\sqrt{\varepsilon}}} - G(z) \right) \tilde{J}^\varepsilon(w)(dz, dt)$$

$$\begin{aligned}
& + \frac{1}{\gamma} \int_{\mathbf{R}^2} \|\Theta_{w,t}^* \mathbf{Q}_{z,\gamma}^\varepsilon - \Theta_{w,t}^* \mathbf{Q}_\gamma\| \tilde{J}^\varepsilon(w)(dz, dt) \\
& + \frac{1}{\gamma} \left\| \int_{\mathbf{R}^2} \Theta_{w,t}^* \mathbf{Q}_\gamma (\Pi_2^* \tilde{J}^\varepsilon(w)(dt) - \Pi_2^* \tilde{J}(w)(dt)) \right\|_1 \\
& \leq C_\gamma \varepsilon^{1/5} J^\varepsilon(w)(\mathbf{R}^2) + C_\gamma \varepsilon^\delta \tilde{J}^\varepsilon(w)(\mathbf{R}^2) \\
& + C_\gamma \varepsilon^{1/10} (1 + \text{Lip } \Theta_{w,t}^* \mathbf{Q}_\gamma),
\end{aligned}$$

where $\text{Lip } \Theta_{w,t}^* \mathbf{Q}_\gamma$ denotes the quantity

$$\text{Lip } \Theta_{w,t}^* \mathbf{Q}_\gamma = \sup_{t \neq t'} \frac{\|\Theta_{w,t}^* \mathbf{Q}_\gamma - \Theta_{w,t'}^* \mathbf{Q}_\gamma\|_1}{1 \wedge |t - t'|^{1/3}}.$$

It is at this stage that the choice (6.37) of distance function becomes clear: with respect to the usual Euclidean distance, the map $t \mapsto \Theta_{w,t}^* \mathbf{Q}_\gamma$ would not be Lipschitz continuous. In this way however, it follows immediately from the Hölder continuity of Brownian motion that $\text{Lip } \Theta_{w,t}^* \mathbf{Q}_\gamma < \infty$.

Furthermore, it follows from the definition of $\Omega_\varepsilon^{(2)}$ that $\tilde{J}^\varepsilon(w)(\mathbf{R}^2)$ is bounded by a constant independent of w and ε . Since $G(x)$ is bounded from below by a constant, this implies that the same is true of $J^\varepsilon(w)(\mathbf{R}^2)$. Combining these bounds, we then obtain

$$\|\mathcal{Q}_\gamma(w, \mathcal{X}) - \mathcal{Q}_\gamma^\varepsilon(w, \mathcal{X})\|_1 \leq C_\gamma \varepsilon^\delta,$$

for some constant C and any $\delta < \frac{1}{16}$.

In order to complete the proof, it thus remains to show that $\inf_{w \in \Omega_\varepsilon} \mathcal{Q}_\gamma^\varepsilon(w, \Omega_\varepsilon^c) \rightarrow 0$ sufficiently fast as $\varepsilon \rightarrow 0$, where Ω_ε^c denotes the complement of Ω_ε . As a consequence of the definition of Ω_ε , this follows if we can show that $\mathbf{Q}_{z,\gamma}^\varepsilon(\Omega_\varepsilon^c) \rightarrow 0$ as $\varepsilon \rightarrow 0$, uniformly over all $z \leq \varepsilon^{-1/3}$. Instead of considering $\mathbf{Q}_{z,\gamma}^\varepsilon$, it is much easier to consider \mathbf{Q}_z^ε , the law of the rescaled random walk (6.1) over the time interval $[0, 1]$.

Observe that $\mathbf{Q}_{z,\gamma}^\varepsilon$ is obtained from \mathbf{Q}_z^ε by first conditioning on the event that γ is reached before the walk becomes negative and then stopping the walk. By Proposition 5.1, the probability of this event is bounded from below by $C\sqrt{\varepsilon}$ for some constant C depending on γ but independent of $z \in [0, \varepsilon^{-1/3}]$. Therefore it is sufficient to find a set Ξ_ε with the following two properties:

- There exists an exponent $\zeta > \frac{1}{2}$ and a constant C such that $\mathbf{Q}_z^\varepsilon(\mathcal{X} \setminus \Xi_\varepsilon) \leq C\varepsilon^\zeta$ for every $\alpha \leq 1$ and every $z \in [0, \varepsilon^{-1/3}]$.
- For every $w \in \Xi_\varepsilon$ and every $t \in [0, 1]$, the path \tilde{w} obtained from w by stopping it at time t belongs to Ω_ε .

In order to determine whether a path w belongs to Ξ_ε , we “coarse-grain” it into pieces of length $\varepsilon^{1/3}$ (the precise exponent is not very important as we do not endeavour to obtain optimal convergence rates) and we impose that the contribution of $\tilde{J}^\varepsilon(w)$ on each piece is very close to what it should be. In other words, setting $I_k \subset [0, 1]$ by $I_k = [\varepsilon^{1/3}k, \varepsilon^{1/3}(k+1))$, we define Ξ_ε as

$$\Xi_\varepsilon = \Omega_\varepsilon^{(1)} \cap \left\{ w : \left| \Pi_2^* \tilde{J}^\varepsilon(w)(I_k) - \frac{a}{2} |I_k| \right| \leq \sqrt{\varepsilon} \right\},$$

where $|I_k|$ denotes the length of the interval I_k .

We first note that if $w \in \Xi_\varepsilon$ and $t \in [0, 1]$, then the path \tilde{w} obtained by stopping w at time t does belong to Ω_ε . Since $\tilde{J}^\varepsilon(\tilde{w}) \leq \tilde{J}^\varepsilon(w)$, $w \in \Omega_\varepsilon^{(1)}$ implies that $\tilde{w} \in \Omega_\varepsilon^{(1)}$. Denote now by η^ε the measure $\eta^\varepsilon = \Pi_2^* \tilde{J}^\varepsilon(\tilde{w})$ and by η the target measure, namely

$\eta = \frac{a}{2} \lambda|_{[0,t]}$, where λ denotes the Lebesgue measure. By the assumption on $\tilde{J}^\varepsilon(w)$, it follows that one has $|\eta^\varepsilon(I_k) - \eta(I_k)| \leq \sqrt{\varepsilon}$ for each of the intervals I_k , except possibly for the interval containing t . As a consequence, denoting by η_k^ε and η_k the restrictions of η^ε and η to I_k , one obtains for all k such that $t \notin I_k$ the bound

$$\|\eta_k^\varepsilon - \eta_k\|_1 \leq \varepsilon^{\frac{1}{9} + \frac{1}{3}} + \varepsilon^{\frac{1}{2}}.$$

(Recall that we use the distance function (6.37), this is the reason for the exponent $\frac{1}{9}$.) Summing over k and using the fact that $\|\eta_k^\varepsilon - \eta_k\|_1 \leq C\varepsilon^{1/3}$ for the value k such that $t \in I_k$, it follows that one has $\|\eta^\varepsilon - \eta\|_1 \leq C\varepsilon^{1/9}$, which indeed implies that $\tilde{w} \in \Omega_\varepsilon$, at least for ε small enough.

It remains to show that $\mathbf{Q}_z^\varepsilon(\mathcal{E} \setminus \Xi_\varepsilon) \leq C\varepsilon^\zeta$ for sufficiently large ζ . Observe first that $\mathbf{Q}_z^\varepsilon(\mathcal{E} \setminus \Omega_\varepsilon^{(1)}) \leq C\varepsilon^p$ for every $p > 0$ thanks to the fact that the distribution ν of the steps of our random walk has exponential tails.

To obtain the required bound on $\mathbf{Q}_z^\varepsilon(\mathcal{E} \setminus \Xi_\varepsilon)$, we use the fact that $\Pi_2^* \tilde{J}^\varepsilon(w)(I_k)$ consists of the sum of $\varepsilon^{-2/3}$ i.i.d. copies of a random variable Y with law given by

$$Y \stackrel{\text{law}}{=} a\varepsilon \int_0^Z e^{a\sqrt{\varepsilon}z} G(z) dz, \quad \mathcal{D}(Z) = \nu,$$

where ν is the one-step distribution of the underlying random walk. Note now that, as a consequence of Proposition 5.3 and the fact that ν admits some exponential moment, one has

$$\mathbf{E}Y = \frac{a}{2}\varepsilon + \mathcal{O}(\varepsilon^{3/2}),$$

for all ε sufficiently small. Similarly, it is straightforward to check that $\mathbf{E}|Y|^p = \mathcal{O}(\varepsilon^p)$ for every $p > 0$. As a consequence, one has for every $p > 0$ the bound

$$\mathbf{E} \left| \Pi_2^* \tilde{J}^\varepsilon(w)(I_k) - \frac{a}{2}|I_k| \right|^p \leq C\varepsilon^{2p/3},$$

for a constant possibly depending on p . It immediately follows that

$$\mathbf{P} \left(\left| \Pi_2^* \tilde{J}^\varepsilon(w)(I_k) - \frac{a}{2}|I_k| \right| > \sqrt{\varepsilon} \right) \leq C\varepsilon^p,$$

for every $p > 0$. Summing over k and combining this with the previously obtained bound on $\mathbf{Q}_z^\varepsilon(\mathcal{E} \setminus \Omega_\varepsilon^{(1)})$, we conclude that $\mathbf{Q}_z^\varepsilon(\mathcal{E} \setminus \Xi_\varepsilon) \leq C\varepsilon^p$ for every $p \geq 0$, which concludes our proof. \square

References

- [AG10] S. ASMUSSEN and P. GLYNN. *Stochastic Simulation: Algorithms and Analysis*. Springer, 2010.
- [AJ06] J. A. ADELL and P. JODRÁ. Exact Kolmogorov and total variation distances between some familiar discrete distributions. *J. Inequal. Appl.* Art. ID 64307, 8.
- [Ald91] D. ALDOUS. The continuum random tree. I. *Ann. Probab.* **19**, no. 1, (1991), 1–28.
- [And75] J. ANDERSON. A random-walk simulation of the Schrödinger equation: H_3^+ . *J. Chem. Phys.* **63**, no. 4, (1975), 1499–1503.
- [Bil99] P. BILLINGSLEY. *Convergence of probability measures*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons Inc., New York, 1999.

- [Buc04] J. BUCKLEW. *Introduction to Rare Event Simulation*. Springer, 2004.
- [CA80] D. CEPERLEY and B. ALDER. Ground state of electron gas by a stochastic method. *Phys. Rev. Lett.* **45**, no. 7, (1980), 566–569.
- [Cri11] D. CRISAN. Discretizing the continuous time filtering problem. order of convergence. In D. CRISAN and B. ROISOVSKII, eds., *The Oxford Handbook of Nonlinear Filtering*, 572–597. Oxford University Press, 2011.
- [Daw93] D. A. DAWSON. Measure-valued Markov processes. In *École d’Été de Probabilités de Saint-Flour XXI—1991*, vol. 1541 of *Lecture Notes in Math.*, 1–260. Springer, Berlin, 1993.
- [DBC98] C. DELLAGO, P. BOLHUIS, and D. CHANDLER. Efficient transition path sampling: Application to lennard-jones cluster rearrangements. *J. Chem. Phys.* **108**, no. 22, (1998), 9236–9245.
- [DD11] T. DEAN and P. DUPUIS. The design and analysis of a generalized RESTART/DPR algorithm for rare event simulation. *Ann. Oper. Res.* **189**, (2011), 63–102.
- [dDG05] N. DE FREITAS, A. DOUCET, and N. E. GORDON. *Sequential Monte Carlo Methods in Practice*. Springer, 2005.
- [EK86] S. N. ETHIER and T. G. KURTZ. *Markov Processes*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons Inc., New York, 1986.
- [FG10] A. FIGALLI and N. GIGLI. A new transportation distance between non-negative measures, with applications to gradients flows with Dirichlet boundary conditions. *J. Math. Pures Appl. (9)* **94**, no. 2, (2010), 107–130.
- [FINR04] L. R. G. FONTES, M. ISOPI, C. M. NEWMAN, and K. RAVISHANKAR. The Brownian web: characterization and convergence. *Ann. Probab.* **32**, no. 4, (2004), 2857–2883.
- [Fra73] D. F. FRASER. The rate of convergence of a random walk to Brownian motion. *Ann. Probability* **1**, (1973), 699–701.
- [FS96] D. FRENKEL and B. SMIT. *Understanding Molecular Simulation*. Academic Press, 1996.
- [GS71] R. GRIMM and R. STORER. Monte-Carlo solution of Schrödinger’s equation. *J. Comp. Phys.* **7**, no. 1, (1971), 134–156.
- [HT99] Z. HARASZTI and J. TOWNSEND. The theory of direct probability redistribution and its application to rare event simulation. *ACM Transactions on Modelling and Computer Simulation* **9**, (1999), 105–140.
- [Hut09] M. HUTZENTHALER. The virgin island model. *Electron. J. Probab.* **14**, (2009), no. 39, 1117–1161.
- [Itô72] K. ITÔ. Poisson point processes attached to Markov processes. In *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability (Univ. California, Berkeley, Calif., 1970/1971)*, Vol. III: *Probability theory*, 225–239. Univ. California Press, Berkeley, Calif., 1972.
- [Jag75] P. JAGERS. *Branching processes with biological applications*. Wiley-Interscience [John Wiley & Sons], London, 1975. Wiley Series in Probability and Mathematical Statistics—Applied Probability and Statistics.
- [Kal62] M. KALOS. Monte Carlo calculations of the ground state of three- and four-body nuclei. *Phys. Rev.* **128**, no. 4, (1962), 1791–1795.
- [KL11] J. KOLORENČ and M. LUBOS. Applications of quantum Monte Carlo methods in condensed systems. *Rep. Prog. Phys.* **74**, (2011), 1–28.
- [LG07] J.-F. LE GALL. The topological structure of scaling limits of large planar maps. *Invent. Math.* **169**, no. 3, (2007), 621–670.

- [LG10] J.-F. LE GALL. Geodesics in large planar maps and in the Brownian map. *Acta Math.* **205**, no. 2, (2010), 287–360.
- [Liu02] J. LIU. *Monte Carlo Strategies in Scientific Computing*. Springer, 2002.
- [LL10] G. F. LAWLER and V. LIMIC. *Random walk: a modern introduction*, vol. 123 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2010.
- [Lor63] E. N. LORENZ. Deterministic nonperiodic flow. *J. Atmos. Sci.* **20**, (1963), 130–141.
- [MGG94] R. N. MILLER, M. GHIL, and F. GAUTHIEZ. Advanced data assimilation in strongly nonlinear dynamical systems. *J. Atmospheric Sci.* **51**, no. 8, (1994), 1037–1056.
- [MTAC12] M. MORZFELD, X. TU, E. ATKINS, and A. J. CHORIN. A random map implementation of implicit filters. *J. Comp. Phys.* **231**, no. 4, (2012), 2049 – 2066.
- [NR02] D. NUALART and A. RASCANU. Differential equations driven by fractional Brownian motion. *Collect. Math.* **53**, no. 1, (2002), 55–81.
- [Pra08] A. PRATELLI. On the sufficiency of c -cyclical monotonicity for optimality of transport plans. *Math. Z.* **258**, no. 3, (2008), 677–690.
- [PY82] J. PITMAN and M. YOR. A decomposition of Bessel bridges. *Z. Wahrsch. Verw. Gebiete* **59**, no. 4, (1982), 425–457.
- [PZ07] S. PESZAT and J. ZABCZYK. *Stochastic partial differential equations with Lévy noise*, vol. 113 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 2007. An evolution equation approach.
- [RW94] L. C. G. ROGERS and D. WILLIAMS. *Diffusions, Markov processes, and martingales. Vol. 1*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons Ltd., Chichester, second ed., 1994. Foundations.
- [RY91] D. REVUZ and M. YOR. *Continuous martingales and Brownian motion*, vol. 293 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1991.
- [SS08] R. SUN and J. M. SWART. The Brownian net. *Ann. Probab.* **36**, no. 3, (2008), 1153–1208.
- [VAVA91] M. VILLEN-ALTAMIRANO and J. VILLEN-ALTAMIRANO. RESTART: A method method for accelerating rare event simulations. *Proc. of the 13th international teletraffic congress, queuing, performance and control in ATM* **9**, (1991), 71–76.
- [VEWa] E. VANDEN-EIJNDEN and J. WEARE. Data assimilation in the low noise, accurate observation regime with application to the kuroshio current. *submitted*.
- [VEWb] E. VANDEN-EIJNDEN and J. WEARE. Rare event simulation for small noise diffusions. *Communications on Pure and Applied Mathematics*, *accepted*.
- [Vil03] C. VILLANI. *Topics in optimal transportation*, vol. 58 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2003.